[SQUEAKING]

[RUSTLING]

[CLICKING]

SCOTT HUGHES: All right. So welcome back. I had a little bit of a break. Before I get into-- I go over my quick recap, you hopefully all have seen the announcements that we're going to delay the due date of the next problem set until Tuesday. That's in part because some of the material that appears on it is what we're going to talk about in today's lecture. So I want to make sure you've at least seen a lecture on all the topics before you try to do some of the problems on it.

In truth, you could probably-- based on things I talked about in the previous lecture, you could probably deduce a lot of what you need to do, but still, it's good to go over it a little bit more methodically.

And I'm trying to think. Is there anything else I want to say? Yeah. So then the next problem set will be posted a week from-- is that right? Yeah, I think we post a week from today.

So let me just get back to what you were talking about last time. So I began introducing some geometric concepts and some quantities that we use to describe matter. We've done a lot of things so far that's appropriate for describing kind of the kinematics of particles, but that's rather restrictive. We want to talk about a broader class of things than that. And I began by introducing something which is, again, fairly simple, but it's a useful tool for beginning to think about how we are going to mathematically categorize certain important types of matter.

Pardon me while I put some of my notes away. Also, pardon me. I'm recovering from a-- I caught a terrible cold over the long weekend, and I'm coughing incessantly. So I will be occasionally sucking on a cough drop.

So I introduced the quantity called the number four-vector. And that is given by-- imagine you have sort of a-- it's best to think about this in the context of something

like dust. So you have some kind of non-interacting little agglomeration of tiny particles. And in the rest frame of an element of this dust-- so imagine you go into-- well, I'm going to define the rest frame of that element by saying imagine you've got a cubic nanometer of it or something like that. And you go into the frame where everything in that cubic nanometer is on average at rest.

We'll call n sub 0 the rest density. So I forgot to write that down. So that's the rest density. The rest-- pardon me-- number density.

So I go into that rest frame. That tells me how many little dust particles there are per cubic volume, per element of volume. Multiply that by the four-vector that-- excuse me-- the four-velocity describing that element. And we're going to call that capital N. So that's a vector that describes-- its components describe the number density in some other frame of reference, in any frame of reference. And the flux, talking about how those dust particles flow from one element to another.

So we spent a little bit of time talking about how to define volume elements in a covariant fashion. I'm not going to go through that, but part of the punchline of doing that is that we're going to define a sort of what I call a covariant formulation of conservation of number. It's going to be the spacetime divergence of that vector is equal to 0. That holds-- this equation holds in all frames of reference.

If I choose a particular frame of reference, in other words, I define a particular time, I define a particular set of spatial coordinates, I can then break this up and say that this is equivalent to saying that the time derivative of the t component of that plus the spatial divergence, the spatial confluence is that they sum to 0. I can also take this thing and integrate over a-- oops-- it would help to have one element defined here.

I integrate this over a four-dimensional volume. And by my conservation law, I must get 0 when I do that. Again, having chosen a particular frame, defined what time means, defined what space means, I can then split this out into a statement that the time derivative of the volume integral of the time component of this thing is balanced by the flux of the number across the boundaries of a particular three-volume.

I remind you this notation, this sort of delta V3 means the boundaries of the V3

that's used under this integral. So we talked a little bit about a few other four-vectors that are used. In particular I introduced some stuff that we use in electricity and magnetism.

But I want to switch gears. In particular, I want to introduce and discuss in some detail one of the most important tensors we are going to use all term. So I motivated this by saying imagine I have a cloud of dust. And I won't recreate one, because I am getting over a cold. And I don't want to breathe chalk dust.

So I imagine I make a little cloud of these things here. So far we've characterized this thing by just counting the number of dust particles that are in there. But that's not all there is to it, right? That dust has other properties. And so the next thing which I would like to do is let's consider the energy and the momentum of every particle of dust in that cloud.

So let's just imagine, again, we're going to start with something simple, and then we'll kind of walk up from there to a more generic situation. Let's imagine that my cloud consists of particles that are all identical, and each dust particle has the same rest mass. So suppose this guy has a rest mass m.

So one of the things that I might want to do is in addition to characterizing this dust by saying that the number density-- [COUGHING] excuse me-- in addition to saying that the number density in a particular frame is n or n0, let's go into the rest frame of this thing. I might be interested in knowing about the rest energy density.

So I'm going to denote the rest energy density by rho sub 0. Each particle has a rest energy of m. Remember, c is equal to 1. So mc squared, if you want to hold to the formula that we've all known and loved since we were babies. So that's the rest energy.

And then how many-- if I want to get the rest energy density, I count up the number in each volume, or the number per unit volume. So that's m times n0. So that's the rest energy density of this thing. Great.

Let's now ask, OK, so that's what it looks like in the rest frame of this dust element. Let's imagine that I now bop into a frame that is moving with some speed relative to this rest frame. In this frame, the energy density-- I'm not going to say rest energy

density, because I'm no longer in the rest frame.

But the energy density in this frame, which I will denote by rho without the subscript 0 is going to be the energy of each particle. So that's gamma times m. And the volume is going to be length-- it's going to be a little bit larger. Sorry. The volume is smaller because of a length contraction. And so this guy actually gets boosted up to gamma and 0.

So in this frame, the energy density is larger than the rest energy-- excuse me-- the rest energy density with a gamma squared factor. Straightforward algebra. But I want you to stop and think about what that's telling us.

If this rho were a component of a four-vector, is there any way I could get a gamma squared out of this thing by bopping between reference frames? No, right? When I do that, it's linear in the Lorentz transformation matrix. There's no way with a linear transformation that I would pick up two powers of gamma.

This is not how a-- you guys are probably all used to from Newtonian physics of thinking of energy density as a scalar. This is not a Lorentz scalar. A scalar is the same in all frames of reference. So this is a different quantity.

So it's neither a four-vector component nor a scalar. It's got two of those things, so it's not going to surprise you that what's actually going on here is we are picking out-- what we've done here is we've picked out a particular component of a tensor.

Let's think a little bit more methodically about what tensor that must be. So when I originally wrote down my rho 0 here, I sort of just argued on physical grounds that it is the energy of every particle times the number of particles per unit volume of that thing. Well, the energy that I use in this thing doing it in both the rest frame and in the other frame, those are the time-like components of a particular four-vector.

So we assembled rho by combining energy, which is the time-like component of the four-momentum with number density, which is the time-like component of the number vector that I just recapped a few minutes ago.

So those are two time-like components of four-vectors here. So that rho is something like-- well, let's put it this way. I can write this as pt nt, which tells me that it belongs to a tensor. So let's define this as some tt of a tensor that I've not

carefully introduced yet. There is some underlying tensor that we have built by looking at the-- there are various names for this. We'll call it the tensor product of two different four-vectors.

So if I write this in the kind of abstract tensor notation that I've used occasionally-- in some of your textbooks this would be written with a bold faced capital T, I will use a double over line for this. We might say that this t is the tensor product of the number vector with the four-momentum.

Now, the number vector is itself just n0 times the four-velocity. And my momentum, since all the particles, I'm assuming, are the same is just the rest mass times the four-momentum. This thing actually looks kind of like the four-velocity times four-velocity with a prefactor And that prefactor is nothing more than the rest density.

So another way to say this is that if I make my cloud of dust, and I go in and I look at every little element in it, I take the rest energy density of every element in that dust. I construct the tensor that comes from making the tensor product of the four-velocity of that dust with itself. That is what this geometric object is equal to. If you want to write this in index notation, which is how we will write it 99.8% of the time, we would say that component alpha beta of this quantity t is that rest density times u alpha u beta.

So this is-- in terms of the physics we're going to do this term, this is perhaps the most important quantity that we're going to talk about. So it's worth understanding what this thing is really telling us. So if I wanted to get these components, the alpha beta, out of this thing from this sort of abstract notation of the tensor, I'll remind you that we can do this by taking the tensor and plugging into its slots the basis 1 forms, which tell me something about-- basis 1 forms are really useful for sort of measuring fluxes across in a particular direction.

And in fact, this quantity has the geometric interpretation-- think of this as the flux of momentum component alpha in the beta direction. So if we look at this component by component, it's worth doing that. So let's do the tt or 00 component.

So this is the one we've already done. This is just rho 0 ut ut. It's just rho in that thing.

According to the words I've written down here, this is the flux of pt in the t direction. So pt tells me about energy. The flux of energy in the energy direction refers to-- so here's my water flowing in the time-like direction, just sitting there apparently doing nothing, but no, it's moving through time.

This is energy flowing through time. It's energy density, all sort of locked up in stable equilibrium here. But if somebody comes in someday with a nice cup of anti-water and combines it, we get to have all that energy released. And we can enjoy that for a few femtoseconds before we evaporate.

Let's look at the other components. So T0i, you can take the definition and plug it in in terms of the components. It's obviously ut times ui, but still use the definition by words. So this is the flux of p sub t in the xi direction.

This is talking about energy moving in a particular direction. This is nothing more than energy flux, energy sort of flux that we are used to think about, not flowing through time, but flowing through space.

Now, T0i, this is the flux of momentum component i in the t direction. So this tells me about momentum density. So if I had this stuff flowing, there's some momentum associated with that flow. You can count up the amount of mass and every little volume. Divide by that volume. Take some ratios. That is the density of momentum.

And finally, this last one, flux of p i in the xj direction, there's really no great wisdom for that. This is nothing more than momentum flux.

If I have a bucket of water and things are sort of sloshing around, there's some momentum moving, and the whole assembly is moving in some direction, we can get a flow of momentum going in kind of a non-normal direction. As long as it goes in a normal direction, it might be moving along its flow. I'm going to talk about that in a few moments when I talk about a few different kinds of stress energy tensors.

The key thing which I want you to be aware of is that under the hood of this thing, we're going to talk about different kinds of stress energy tensors for a little later in this class.

But this basic interpretation of the way to think about the different components, it

holds for all of them. So this is good intuition to have.

One thing which is worth noting is if you look at-- if you actually-- let's take the form of the stress energy tensor that corresponds to the dust. So I'm going to write out all four of my components. I've already given you T00. T0i, I can write that as gamma squared rho 0 vi T0i equals gamma squared rho 0 vi.

This guy is gamma squared rho 0 vi vj. There's two elements of symmetry here, which I want to emphasize. Notice energy density and momentum flux. Sorry. Energy flux and momentum density, those words are important. Energy flux and momentum density, they're exactly the same.

That is actually true in all physics. It's clouded by the fact, though, that in most units that we measure things, we don't typically use speed of light equal to 1. So when you look at these kind of quantities in the physics that you are more used to, you would find your energy, your energy flux, your momentum density will have different units, because you'll have factors of c that come into there that convert from 1 momentum, and energy, and time flow, and space flow.

Also, you don't usually include rest mass and rest energy in the energy flows and momentum densities that you've probably been familiar with in the past. So it's when we do relativity, and we set the speed of light equal to 1 that this symmetry between energy flux and momentum density becomes apparent.

Question.

AUDIENCE:	The second one, do you mean rho 1 will be there?
SCOTT HUGHES:	Sorry. This one?
AUDIENCE:	Yeah, on the right-hand side.
SCOTT HUGHES:	On the right-hand side. This is--
AUDIENCE:	Rho i.
SCOTT	No, this is definitely what I mean to write down here. So this is-- what I've done, I've

HUGHES: taken advantage of the fact that I can write-- let's see. Do I have it down somewhere? Yeah. Yeah.

So what I've done over here is I have written the fact-- so u sub t, ut is just the gamma factor. And I didn't explicitly write out what t-- the i components, but you get a similar factor that way. So no, that is correct that way.

AUDIENCE: So that's why you mean the first and the second are the same.

SCOTT HUGHES: Yeah. Yeah. And I just want to emphasize-- I mean, the key thing which I want to emphasize is when you do it in the physics that you guys have generally all seen up till now, it won't look that way. And there's two causes of that. One is you're not used to necessarily working in units where the speed of light is 1. And two, rest energy is built into these definitions. That shifts things a little bit.

The other thing which I want to emphasize is notice this is symmetric under exchange of indices. Now, it's obvious for dust. It comes right back to the fact that these guys-- that those indices enter like so. It's an exterior product of the four-velocity with itself.

I'm not going to go through in detail, but there is a physical motivation that I will sketch in a few moments that argues that it must be symmetric like this in all physical cases. I have a detailed sketch of this. I keep using the word sketch, because my mind wants to say proof, but it ain't a proof. It's sort of a motivation. And so I will post it for everyone to go through, but I will sort of illuminate it in a couple of moments.

Basically, if this is not true, if it's not the case this is symmetric, a physical absurdity can be set up. I'll describe that in a few moments. So most of the world is not dust. So this is a decent example for helping to understand things. But we are going to want to do more complicated and interesting examples.

Pardon me. Just one second. Let me just check here with this page. Yeah.

So for us, the way that we are going to do this, there is a fairly general recipe that one can imagine applying to this. I'm going to save it for a little later in the class. It sort of borrows some techniques from field theory.

Basically, if you can write down a Lagrangian density that describes the fifth system that you're under study, there's a particular variation you can do, that stress energy tensor emerges from. But for now, the key thing which I want to say is that we basically are going to deduce what the stress energy tensor looks like by essentially going into a particular frame-- we'll call it the rest frame. It's usually the rest frame of some element of the material being studied-- and thinking carefully about this physical definition of what the different components mean. This is not the most rigorous way to do it, but it's a good way to get started and develop some intuition.

So let me give one example that for many of us in astrophysics, this is probably the one stress energy tensor that we write down and use over, and over, and over again in our career. And it's rare we do anything more than this in many, many cases. This is called a perfect fluid.

So what is perfect about a perfect fluid? It begs the question here. What is perfect referring to here? So a perfect fluid is a fluid in which there is no energy flow in what I will call the "rest" frame. I put rest in quotes, because you have to sort of define it from the context of what your fluid is here. Basically, what it means is I can find a frame in which each fluid element there is no energy flow there. If a frame exists where that happens, there's a candidate to be a perfect fluid.

And I also require there to be no lateral stresses. Lateral stresses refer to this Tij when i and j are not equal to one another. So this sort of refers to-- imagine that I have some-- let's say going into the board is the y direction.

So if there's like some y stress that is somehow being transported in the x direction, that would be a lateral stress. Physically, this is actually-- that kind of a stress is hugely important when one is studying fluids. And one characterizes it by a quantity known as the viscosity. Viscosity tells me about how stress gets transported-- how momentum, rather, gets transported in a non-normal direction, against the direction in which the fluid is moving.

So my perfect fluid has no viscosity. And I'll just conclude. I want to make-- there's a very important point here. A fluid that has no viscosity is a fluid that doesn't get anything wet. So this refers to when you pour water on yourself, the reason your hand gets wet is that there's some viscosity that actually prevents it-- causes there

to be sort of a shearing force, which causes the water to stick to your skin.

So a perfect fluid, this has been described as the physics of dry water. So it is-- anyone who's in applied math will know. They'll sort of roll their eyes and say, OK, fine, we'll do this sort of infant version of fluid first. And then a lot of the action and a lot of the fun comes from putting in viscosity and doing the real fluids.

For the purposes of our class, what this boils down to is this tells me that the physics of this quantity is totally dominated by the fluid's energy density and its pressure. The pressure is an isotropic spatial stress.

So in this particular frame, where I have no energy being transported, then the stress energy tensor can be represented as energy density as usual up in the-- whoops, over there now-- up in the tt component. There is no energy flux. By symmetry, if there's no energy flux, there is no momentum of density. And my spatial stresses are totally isotropic, and none of them are lateral. So it just looks like this, the diagonal of rho p, p, p.

As I said, this is actually something that we're going to use over, and over, and over again. This is actually a-- so you can actually consider my dust stress energy tensor to be a perfect fluid with no pressure. So this actually subsumes this other one.

I will-- just illustrative purposes, I'll show an example of a case that cannot be thought of as a perfect fluid. But we will tend to use this a lot in our class. And as I will demonstrate in just a moment, we're going to find that this ends up playing an important role in generating gravitational fields. What's interesting about this is you guys are probably used to the idea that mass generates gravity, and then throw in a c square that tells you energy to generating gravity. But it's also-- we're going to see pressure generates gravity. And it's connected to the mathematical structure of this guy here.

So notice I've written an equal dot here. So this is just the way this is represented in this particular frame. I would like to write this in a more covariant form, something that does not rely on me going into a particular frame of reference. So the trick which I'm going to use for this is when I think about that form there-- so the rho piece of it, clearly, what I'm doing there is I'm picking out-- that is-- can be thought of-- as the energy density multiplied by the tensor product of the four-velocity that

describes a particular element of this fluid.

So I can write, again, using this sort of abstract notation. This piece of it looks like this. And so if I go into the frame, if I go into the rest frame of the fluid, that's just my u is 1 and 0. And the spatial components, that builds my upper left-hand corner of this tensor.

How do I get the rest of this? Well, to get the rest of it, these are all sort of picked out of components of the tensor that are orthogonal to u. And I put an Easter egg in p set 1.

You guys developed a tensor that allows me to build-- it's a geometric object that allows me to describe things that are orthogonal to a given four-vector. So if I take the projection tensor that you guys built one piece at 1, which looks like the metric plus the extern-- the tensor product of a four-vector with itself, that gives me the p's that go into that component.

Or if I write this in index notation, I can do it in two ways. I will also emphasize that there are few moments in this class where I sort of urge you to take the sort of long-term memory synapses and switch them on. This is one of those moments.

In a couple of lectures, we're going to introduce the principle of equivalence, which is the physical principle by which we're going to argue how we go from formulas that work in special relativity to formulas that work in general relativity. And by invoking the principle of equivalence, we're going to see that when we want to describe perfect fluids in a general spacetime, not just in special relativity and general relativity, it's exactly this formula. I just need to modify what the metric means, but that will allow me to carry that over.

Let me see. So my notes are a little bit disorganized here. There's a piece that I-- every year I want to clean this bit up, and every year at this time of year, I have 70 gajillion administrative tasks. And I end up getting behind schedule. So I will clean this up on the fly.

So there's a couple of points which I want to make about this. Let me do this point first.

So by virtue of taking a graduate physics class, I'm confident you guys know about Newtonian gravity. One can write the field equation for Newtonian gravity as essentially a differential equation for the potential that governs Newtonian gravitational interaction. So let's call phi sub g the Newtonian gravitational interaction. Whoops. And I can write a field equation that's governing it as essentially its Poisson's equation.

So the Laplace operator acting on that potential is up to a constant equal to the-- we usually learn it in terms of mass density. We're working in units where the speed of light is equal to 1. So it could just as well be the energy density. So you've all kind of seen that.

Now, if we think about how we're going to carry this forward and make gravity a relativistic interaction, this equation should right away make us suspicious, because we spent several minutes earlier today talking about the fact that this is not a scalar. This is a component of a tensor.

A physical theory which tries to pick out just one component of a geometric tensor is not a healthy theory. It would be like if you had sort of learned in E&M that there was a preferred direction of the electric field. If there's a particular set of Maxwell's equations for Ex, and a different set of Maxwell's equations for Ey. That would just be nonsense.

Nature doesn't pick out spatial components of any one being particularly having some weight over the others, and the same thing holds in relativity for spacetime. So when we make this into a relativistic theory, we're going to say if this component of a tensor plays a large role in gravity to make this into a geometric object that makes a large role in gravity, I'm going to have to-- so let's call this my Newtonian equation.

My Einsteinian equation-- I'm going to put an equal sign in quotes here, because there's a lot of details to fill in. But what's going to go on the right-hand side of this has to be something that involves the stress energy tensor.

Newton picks out one component stress energy tensor. Relativity doesn't let me pick out particular components. So whatever I get when I do Einstein's gravity, the whole stress energy tensor is going to be important in setting the source of my

gravity.

That then sort of says, well, then what the heck do you do with this left-hand side? That is, in fact, going to be starting probably on Tuesday the subject of the next couple of lectures, basically going all the way up to spring break. The week before spring break is when we complete the story of what goes on the right-hand side-- excuse me-- the left-hand side of this equation.

But what I will tell you is that it is indeed going to involve two derivatives of a potential-like object. And the potential-like object is actually going to turn out to be the metric of spacetime. So that's kind of where we're going with this.

Let's do a little bit more physics with the stress energy tensor. So I have somewhat more detailed notes, which I will post online, which I am not going to go through in great detail here. But I'm going to kind of sketch this.

So one is that I would like to prove-- again, that word is a little bit of an overstatement, but at least motivate the symmetry of this tensor.

I'm just going to focus on the spatial bits of this. That'll be enough. Like I said, you can kind of see that T0i and ti0 are the same thing by thinking about the physical meaning of four-momentum and what a flux of four-momentum is.

This one, there's kind of a cute calculation you can do. So imagine you have some little cube of stuff. It could be immersed in some field or fluid, something that is described by a field of stress energy.

And so to start out with, let's look at how the flux-- remember what t alpha beta, or really Tij tells me about is the flux of momentum in a particular direction. It's telling me about the amount of momentum going into this box on one side and coming out on the other.

So I'm going to look at the momentum flux into and out of this box. And so what I'm going to do is let's look at the momentum going into the sides that are-- so let's call this the top and the bottom here.

What I'm going to do then is number the four sides, the other four sides of the box, not the top and the bottom-- so let's call this side, which is sort of facing away from

us here, that side one. The one that is facing us that sort of points out in the x direction, I'll call that two. This side here, I will call three. And the one that is on the back I will call four. Apologies for the little bit of a busy picture here. But just what you want to do is sort of imagine there's a cube in front of you, and you go around and label the four sides.

What we'd like to do is calculate what is the force that is flowing through each of these four sides, 1, 2, 3, and 4. Just ignore the top and the bottom for just a second. I'll describe why that is in just a moment.

So if I look at the force on face one, well, face one, it is-- pardon me-- I mislabeled my sides. I want to be synced up with my notes, and so I realize it's annoying. And I'm very sorry about that. But if I mess this up, I will get out of sync with what I have written down.

So the one that is facing us is side one. Two is on the right-hand side. Three is the back. Four is on the left-hand side.

My apologies for that. But I think it's important we get that right so I don't get out of sync with myself.

So the total force on face one is what I get by basically adding up all the flux of momentum flowing through face one. So force on face one, that is-- the i-th component of that is what I get when I integrate that over face one, which is normal to the x-axis. Each side of the cube is little l. So it's Tix l squared.

Force on face two, now this is the one that is normal for the y-axis. And so this guy looks like Tiy l squared. And you continue this F3-- if you look at, and you if you assume that this thing is small, it's approximately the same as the force on F1. It'll become equal in the limit of the q becoming infinitesimally small. And this is approximately the same as the force on face two.

So that tells us that there's no unbalanced force on this thing, which is great. Basically, it means that there's no unbalanced force causing this little element to accelerate away.

One of the reasons why I am focusing on these four sides, though, is I would also like

to consider torques that are acting on this thing.

And here's where I'm going to skip over a couple of details and just leave a few notes for you guys to look at, because we're a little short on time with other things. This calculation is straightforward, but it gets-- I already screwed up a detail here. I don't want to sort of risk messing up a few other things.

What I want to do is consider the torques about an axis that sort of runs through the middle of this thing that goes from the top and the bottom of this. So go through and just add up all the torques associated with these little forces about an axis in the center of-- that goes through the center of this cube.

So I'm going to-- like I said, I'm going to leave out the details, but basically, you go through, and you do the usual r cross f to get the forces, the torques associated with each of these different faces. And what you'll find when you do this is that there is a net torque that looks like l cubed times Txy minus Tyx. Again, though, this scales with the size of the cube.

So you look at this and kind of go, well, who cares? It's a little bit there. Maybe it spins up a little bit. But in the limit, the thing going to 0, there's no net effect.

Well, let's be a little bit careful about this. What is the moment of inertia of this cube? I don't know exactly, but I know that it's going to be something like the mass of this cube. We'll call that l cubed times its mass density, or its energy density. And it's going to involve two powers of the only length scale characterizing this thing.

And there will be some prefactor alpha. There will be some number that's related to the geometry of this. An integral or two will easily work out what the alpha is and make this more precise. But the key bit which I want to emphasize is that this is something that scales as the size of the fifth power.

So yeah, the torque vanishes as l goes to 0, but it does so with the cubed power. It doesn't vanish as quickly as the moment of inertia vanishes. And I'll remind you, the angular acceleration of the cube, theta double dot, is the torque divided by the moment of inertia. So this is something that is proportional to Txy minus Tyx divided by l squared.

So yeah, the torque does vanish as l goes to 0, but the moment of inertia vanishes

more rapidly as l goes to 0. That sort of suggests that if we're in a screw universe, little microscopic vortices are just going to start randomly oscillating in any cup of water that you pour in front of you.

Now, I don't have a proof that nature abhors that, but it seems pretty screwy. And so the case that most textbooks make at this point is to say, physics indicates we must have Txy strictly equal to Tyx in all cases, independent of what the size of this thing is. That had better just be a bloody 0 in the numerator in order to prevent this physical absurdity from being set up.

Repeat the exercise by looking at torques around the other axes. And that drives you to the statement that this thing must, in general, be spatially symmetric. And again, physics of the way that energy and momentum behave in relativity makes our sort of timespace component symmetric as well.

I emphasize this is not a proof. This is really just a physical motivation. This is the kind of thing that I'm a bloody astrophysicist. I like this kind of stuff.

There is a different way of developing the stress energy tensor, as I said, that comes from a variational principle sort of based on sort of quantum field theoretic type of methods. And the symmetry in that case is manifest. It really does sort of come up.

But this is a good way of just motivating the fact that you're not going to have any non-diagonal-- or excuse-- non-symmetric stress energy tensors. If you did, you would have really bizarre matter.

Now, the stress energy tensor has this physical interpretation that it tells me about the flow of energy and the flow of momentum in spacetime. As such, it is the tool by which we are going to put conservation of energy and conservation of momentum into our theory.

And the way we're going to do this is with a remarkably simple equation. The spacetime divergence of t alpha beta-- pardon me-- must equal 0. This is a covariant formulation of both conservation of energy and conservation of momentum.

And if you want to say, well, which one? Is it energy or is it momentum? You can't

say that in general. I can only say that after I have picked a particular reference frame, because it's only once I have defined time and I've defined space that I've actually defined energy and momentum. Prior to choosing a particular reference frame, all I have is four-momentum. One observer's energy is another observer's superposition of energy and momentum.

Once I have picked a particular frame-- so once I have picked a particular frame, then if I evaluate d alpha T alpha either 0 or t, this is what conservation of energy looks like in that frame. So in that frame, here is conservation of energy.

I can fit it.

Hang on a second. No, no, no. I had it right the first time. Yeah. Sorry. I was just trying to make sure I got my indices lined up properly. So apologies you can't quite read the bottom one so well.

The top one is conservation of energy in that particular frame. The second one is conservation of momentum in that particular frame.

And again, the key thing which I want to emphasize is the covariant statement basically puts both of them together into a single equation. I can only sensibly state conservation of energy and conservation of momentum according to some particular observer.

Now, I can repeat the game that I had done before with the number vector and turned these conservation laws into integral equations as well. Let me do it for energy.

So if I take that integral equation and integrate over a three-volume, or if you prefer, I can take the original covariant formulation and integrate that over a four-volume. With a little bit of manipulation akin to the way I manipulated the conserv-- the integrals associated with the number vector in the previous lecture, you can write down a law that looks like this.

And so this-- again, what I've done here is I've chosen a particular set of timeline coordinates. And the language that we often use in relativity is we basically say I'm going to take a single slice of time. And I would say that the rate of change of energy-- so integrate energy over a volume. It's the total energy in that V3. The rate

of change of that thing is balanced by the amount of energy flowing into or out of through the boundaries of that volume.

Add an extra index here. So make this-- make one of these be a j. Make this guy be a j. And you've got a similar statement for the conservation of momentum.

This is a particular trick. I have a set of typed up notes that just sort of clean this up a little bit. They're basically exactly this point, but I'll put them online after today's lecture. You're going to want to use this on one of the p set problems this week.

So you guys are going to do a couple exercises that involve integrating and actually finding essentially moments of the left-hand side of this thing. You're going to do a few exercise where you take advantage of this formulation of conservation of energy and momentum to derive a few identities involved in the stress energy tensor, one of which is another Easter egg that we're going to use quite a bit in a future lecture.

So let me just wrap up our discussion of the stress energy tensor by just doing two more examples. And then I'm going to sort of begin switching gears a little bit.

So far I have only described-- essentially, I've really only described perfect fluids. Dust is a perfect fluid with no pressure. There are a lot of other kinds of materials that we want to work with in the universe.

One of which-- it's unlikely many of you are going to use this, but it's actually the stress energy tensor on which I have built-- well, I guess Alex is going to use it a little bit-- I have built a big chunk of my career. Suppose you have-- we kind of talked a little bit about how the four-velocity and the four-momentum are really only good for talking about like the kinematics of particles.

Well, actually, particles aren't a bad thing to focus on in some of your studies. And much of my research is actually based on the idea of thinking about a binary system as one member being a black hole, and its companion being a particle-like object that is a particular limit of it, a particle-like object that orbits it.

So a really simple stress energy tensor-- and I'm just throwing it out here, because I think it nicely illustrates the principle. A point particle with rest mass, we'll call it

m0. And I'm going to say it's moving on a particular world line through spacetime.

So the way we define a world line is we just say it's some four-vector that describes the place of this thing from some chosen origin. And it's generally most convenient to parameterize it by the proper time of whatever object or creature is moving on that world line. And so it's a lot like dust, only there's no volume. It's sort of like one particle of dust.

And so the stress energy tensor we use for this-- let me back up for just a second here. When you guys learned about electricity and magnetism, one of the first things you learn about are point charges. And then a little bit later you learn about charge distributions. And you have charge densities.

And then, usually, at some point in a class-- it's often like junior level E&M-- we say, how do you describe the density of a point charge? That's where you learn about the Dirac delta function.

Well, if I have a point particle, I'm going to need to describe this thing's energy density as essentially a Dirac delta function. And so what we do is imagining that these things-- so this is the four-velocity of that body. They might be functions of time as this guy is moving along here.

What we do is we introduce a Dirac delta function as this thing moves along through spacetime. What this does-- you can check the dimensions. This gives me exactly what we need in order to have something that's dimensionally correct, and has all the symmetries and all the properties that describe a particle moving at a particular four-velocity through spacetime.

Now, you might want to just-- it's the thing which I kind of want to pause on for a second. You go, what the hell do you do with this, right? That's kind of inconvenient. Well, the trick we use to sort of clean up that Dirac delta function, it's very much like what you do when you encounter multidimensional delta functions in basic physics. You just build it out of a bunch of one-dimensional delta functions.

Likewise, you'll have a term with a y component and the z component. Then you use the rule that if I integrate a function of x against a delta function whose argument is itself a function of x, I evaluate that function at the 0's of g. So let's say x0 is where

g has a 0. Normalizing to the first derivative of g evaluated at that 0.

When you put all of that together, that basically means you can do this somewhat abstract integral formally. I mean, exactly. You can just do it analytically.

And what you do is you choose one of the delta functions to apply it to. Traditionally, people apply it to the timeline component. And what is the derivative of the world line z component with respect to proper time? It's the 0 component of the four-velocity. So you just divide by the 0 component of the four-velocity, and then you're left with a three-dimensional delta function for the spatial trajectory of this thing through all space.

So this is an example of one that just, again, kind of the intuition is one of the things I want to emphasize. Notice we had-- this has the symmetries that we wanted to have.

It's not hard to show that this. You can think of this as essentially being kind like a gamma factor. This all ends up giving me just what we need for this thing to have the right transformation properties. And it does, in fact, play a role in some-- well, I'll say in some research that's near and dear to my heart.

Let me do another example. Suppose you want to know the stress energy tensor associated with a given electric and magnetic field. Well, first, let me just quote for you the exact answer, which is most compactly written if we use that Faraday tensor F, which describes electromagnetic fields in a frame-independent fashion, the way that I introduced it in the last lecture.

So in units where basically everything but pi is set equal to 1, it ends up turning into this. So that's a bit of a mouthful. Let's go in and look at particular components of it, though.

So let's say I go into a particular frame. I fill in my Faraday tensor with the form of electromagnetic field that I introduced last time. And I'll just go through, and I'll evaluate all the different components of this thing.

So what you find when you fill this in is your T00 component, 1 over 8 pi e squared plus b squared. That's good. Hopefully, y'all remember from basic E&M, the energy density of an E-field is e squared over 8 pi and the right system of units. Energy

density of a b field is b squared over 8 pi if you work in God's units.

Let's do the timespace component. So again, hack through that mess there. This is going to be something that is a vector. In fact, it's the Poynting vector. Could it be anything else? If you use the recipe that I sort of suggested as the easiest way to approach this, this is kind of what you would have guessed for something like that.

The bit that's actually kind of hard is then trying to get the spatial stresses of this thing. And here I don't have any great intuition for this one to convey to you. It's derived in textbooks like Griffiths and Jackson. So I'll just quote for you the result.

So you get one term that looks like E squared plus b squared on the diagonals. Then there's a correction, which looks like this.

I want to just quickly call out one example so you can see what the significance of this example is. So suppose you have something like a pair of capacitors. And there's just a uniform electric field between them. You want to evaluate the stress energy tensor between those pairs of capacitors.

So my spatial electric field, let's say it just points in the Ex direction. And it's constant. So when you actually evaluate this guy, there's no energy flow. There's no magnetic field. So there's no Poynting vector. You, of course, have E squared over 8 pi for your energy density. Very different from a perfect fluid.

And this kind of makes sense. That's sort of telling you that there is a stress that if you have a pair of plain parallel capacitors, you tend to attract the plates to each other.

But there is a pressure associated with that electric field that actually goes in the other directions. This is your x direction. This is y and z. And as a consequence of this-- so there's some stuff which we're not going to do too much with in this class, but I may give you some pointers on this.

Electric and magnetic fields, they generate pressures, at least in certain directions. They kind of generate like an anisotropic pressure. And when we start coupling this to gravity, you can get electric fields and magnetic fields that contribute non-negligibly to the gravity of their object.

That is it for the stress energy tensor. As I said, we are going to use this guy over, and over, and over again. And the reason for this does come back to that little motivation that I gave probably about 45 minutes ago where we sort of looked at the Newtonian field equation, and then said, picking out a particular scalar as the source of gravity makes no sense in a relativistic covariant theory. It's got to be the whole tensor. And indeed, this is the one-- well, not the E&M one, but the general notion of a stress energy tensor is the one that we are going to use for that.

So we have about 10 minutes left. And so I would like to start the process of switching gears. Before I do that, are there any questions? I will clean the board.

So the reason we are switching gears is we now have probably the most important physical tools that we need in order to start thinking about making a relativistic theory of gravity. But we need a few more mathematical tools. In particular, I'm going to argue in a couple of lectures that flat spacetime is not sufficient for us to build a theory of gravity. We're going to need to-- first of all, you're going to have to understand what that means. And we're not quite ready to go there. So for now, it's just fancy words.

But I'm going to have to introduce some kind of a notion of curvature into things. What does that even mean really? We need to have the tools to do that. And so as a prelude to going in that direction-- I will call this my prelude to curvature-- what we're going to start doing is flat spacetime in curvilinear coordinates.

And the importance of doing this, why this is going to be useful to us is that it will introduce a-- it'll keep the physics simple. It's still going to be special relativity, but it's now going to special relativity using a mathematical structure in which the basis vectors are no longer constant. So that's going to allow us to begin making a couple of the mathematical tools that are necessary to build gravity into this theory.

So we'll start by replacing this with just simple plain polar coordinates mapped in the usual way. So if we want to go back and forth, well, at least one way transformation, I build x and y from r and phi in that usual way. There's an inverse mapping as well, which involves trig functions. So I'm not going to write it down.

One point which I really want to emphasize here is we are going to continue to use a

coordinate basis. And to remind you what that means, so a coordinate basis means that the differential displacement vector from an event a to a nearby event b is simply related by differentials of my coordinate contracted onto my basis vectors.

But when I'm working in a curvilinear coordinate system like this, that means one of them has a somewhat different form from what you are used to. When you guys talk about the differential of a displacement in just about every physics class up to now, if you have a differential angle, you'd want to throw an r. So this has the dimensions of length. We ain't going to do that.

And so what this means, since this is an angle, every component of this has the dimensions of length. That means that my basis vector is going to be something that has the dimensions of length associated with it. This in turn means that this is not going to be a normal basis.

That's unfortunately a somewhat loaded word. What I mean by that is that it has not been normalized. But given that you guys have spent all of your career thinking about the dot product of a unit vector with itself-- of a basis vector with itself being equal to 1, the other meaning of normal might be good for you, too.

The key thing which I want to emphasize here is e phi dot e phi does not equal 1.

One other little bit of notation which I would like to introduce-- so we are going to want to talk about transformations between different representations. We've done this-- so far, we have generally focused on moving between different reference frames. I want to generalize this notion. And I'm going to tweak my notation a little bit to indicate the difference.

So I'm going to call capital L alpha mu bar what I get when I just look at the variation of the alpha coordinate system-- sorry-- the unbarred coordinate system with the barred one. So this is just-- there's no deep things here. I just want you to be familiar with the notation which I'm going to use.

I will always reserve lambda for the Lorentz Transformation. So in the interest of time, I'll just write down one of these. So this sort of means like, for instance, suppose that I'm transforming. Let's let barred indicate my polar coordinates. Unbarred be Cartesian.

These things are surprisingly sticky sometimes.

So if I want to transform in one direction between my barred and my unbarred ones, this matrix will go through it. And it looks like this.

So this is the thing which I will call l alpha mu bar. In my notes, I also give the inverse transformation. And I will actually write-- you know what? I have minute. I will write that one down as well.

If you want to know how to go in the other direction, it's just the matrix that inverts this. The reason I decided to take the extra 30 seconds or so to write this down is I want to call out the fact that in this kind of transformation, because of the fact that this is a coordinate basis that has this somewhat unusual property, different elements of the matrix have different dimensions. It's a feature, not a bug.

So this is going to show up. I'm going to do a couple more calculations with this a little bit tomorrow-- sorry-- not tomorrow, on Tuesday. This is going to show up in the way the metric looks. The metric is going to have a very different character in this coordinate system.

And we're going to see the way in which-- it basically boils down to it's going to pick up a non-trivial functional form. It's not going to just be a constant. And that fundamentally reflects the fact that the basis vectors are not constant anymore.

We'll end it there.