[SQUEAKING]

[RUSTLING]

[CLICKING]

SCOTT HUGHES: All right, good afternoon. So today's lecture is a particularly important one. We are going to introduce the main physical principle that underlies general relativity. We'll be spending a bunch of time connecting that principle to the mathematics as the rest of the term unwinds. But today is where we're going to really lay out where the physics is in what is known as the principle of equivalence.

So before I get into that, just a quick recap. So in terms of technical stuff we did last time, the most important thing we did was to introduce these mathematical objects called Christoffel symbols. Christoffel symbols are those capital gammas. We began by thinking of them as just what you get when you look at the derivative of your basis objects. Pardon me a second. There was a lot of chalk on the chalkboard, and I think I inhaled half of it while I was cleaning it.

So take derivatives of your basis vectors. And there, you get something that's linearly related to your basis objects. And the gamma is the set of mathematical quantities that sets that linear relationship between the derivatives and the objects themselves.

One can build under the demand-- so recall how we did this. We have a physical argument that tells us that a quantity we call the covariant derivative, which is a way of using these Christoffel symbols in order to get tensorial derivatives.

So if I look at just the partial derivative of a tensor, if I'm working in a general curvilinear coordinate system, the set of partial derivatives do not constitute components of a tensor. OK? I posted some notes online, demonstrating explicitly why that is. But in a nutshell, the main reason is that, when you try to work this out or when you go and you work this out, you will find that there are additional terms. It was like the derivative of the transformation matrix that come in and spoil the tensoriality of this relationship.

And what the Christoffel symbol does, as you guys are going to prove on P set three, it introduces terms that are exactly the same but with the opposite sign and cancel those things out. So that the covariant derivative is what I get when I correct each index with an appropriate factor of the Christoffel symbol. It is something that is a component of a tensor. In other words, it obeys the transformation laws that tensors must have.

We have a physical argument that essentially is something that is tensorial, in other words, both sides of the equations are good tensor quantities in one representation, have to be tensorial in any other representation. And from that, we deduced that the covariant derivative of the metric must always be 0. And by taking appropriate covariant derivatives of the metric, sort of doing a bit of gymnastics with indices and sort of wiggling things around a little bit, we found that the Christoffel symbol can itself be built out of partial derivatives of the metric. And the result looks like this.

So that's a little bit more of an in-depth recap than I typically do. But I really want to hammer that home because this is an important point. We then began switching gears. And we're going to do something.

So I introduced, at the end of the previous lecture, a fairly silly argument, which, nonetheless, has an important physical concept to it. So where we were going with was the fact that, when we have gravity, we cannot cover all of spacetime with inertial frames, or I should say-- I should watch my wording here-- with an inertial frame. OK? We're, in fact, going to see that you can have a sequence of them. And we're going to find ways of linking them up.

But for now, if you think of special relativity as the theory of physics in which we can have Lorentz coordinates that cover the entire universe, that's out the window. OK, and so I'm going to skip my silly derivation. Part one of it is essentially just a motivation that a gravitational redshift exists.

So I gave a silly little demonstration where you imagined dropping a rock off of a tower. You had some kind of a magical device that converts it into a photon that shoots it back up and then converts it back into a rock. And you showed that, as it climbed out of the gravitational potential, it had to lose some of its energy. Or else, you would get more energy at the top than you put in when you initially dropped it. It's a way of preventing you from making perpetual motion machines.

That is, without a doubt, a flaky argument. It can be made a little bit more rigorous. But I didn't want to. And I will justify the reason I can do this is that this is, in fact, a highly tested experimental fact. Every single one of you has probably used this if you've ever used GPS to find your way to some location where you're supposed to meet a friend.

And what this basically tells me is that, if I imagine a tower of height h, if I shoot light with frequency at the bottom, omega b, when it reaches the top, it'll have a different frequency, omega t, where these are related by omega t equals, and I'm going to put my c squareds back in there, just to keep it all complete. OK?

So this is, I emphasize, a very precisely calibrated experimental result. There are many ways of doing it, some silly, some serious. The silly one was sort of fun. But it does actually capture the experimental fact. OK? The key experimental fact is that it's necessary for light to lose energy as it climbs out of the gravitational field, in essence, because it's possible, thanks to quantum field theory, to convert photons into particles and particles into photons. This is a way of conserving energy.

OK, so this was just part one of this argument that we cannot cover all of spacetime with an inertial frame. So here's part two of the argument. Suppose there was, in fact, a very large region of space-- now let's say it's near the Earth's surface-- that could be covered by a single Lorentz frame. OK?

So I want you to imagine that you're sending your beam of light up. And we're going to take advantage of the fact that light is wave like. So what I want you to do is think of the spacetime trajectory of successive crests of this wave as it climbs out of the gravitational potential.

So I"m going to make my spacetime diagram the way we like to do it with time running up. And so this direction represents height. So down here is the bottom of the tower. So let's just make the drawing clearer. Let's let this tick be the bottom of the tower. This tick over here will be the top.

Now if we were just in special relativity, there was no gravity, we know light would just move on a 45 degree angle on this thing. We don't know yet what gravity is going to do to this light. OK? But you can imagine that whatever it's going to do, it's going to bend it away in some way from the trajectory that special relativity would predict.

So let's just imagine that crest one of this light, it follows some trajectory in spacetime that maybe looks like this. So here is the world line of crest one. So what about the world line of credit two? So let's think about what the world line of credit two must be under this assumption that we can, in fact, cover all of spacetime with a single Lorentz frame. OK?

If that is the case, so if that is the case, no position and no moment is special. OK? Everything is actually translation invariant in both time and space. So the second crest is going to be emitted one wave period later. But there's absolutely nothing special about spacetime one wave period later. And as it moves along, there can't be anything special about this.

If it is to be a global Lorentz frame, whatever the trajectory is that the crest two follows through spacetime, it's got to be congruent with the trajectory of crest one. I've got to be able to just simply slide them on top of each other. OK? That is what this assumption demands.

So this is going to look something like this. OK? Artwork a little bit off. The key thing which I want to emphasize, though, is that if this guy is congruent, the spacing between crest one and crest two at the bottom, let's call it delta tb, there'll be some spacing between the two at the top. We'll call it delta tt. They must be the same.

But that's [INAUDIBLE]. That is just a period of the wave. And that's just up to a factor of 2 pi. That's 1 over the frequency. So this implies delta tb is not equal to delta t at the top. So a frequency at the top is smaller. So the period will actually be a little bit larger.

So when you think about this argument carefully, the weak point is this assumption that I can actually cover the region with a large Lorentz frame. Remember what this is telling me is that, essentially, there is no special direction, right? I can do translation, variance in time and space, and they're all the same. But that's actually completely contradicted by common sense. If there were no special direction, why did it go down? There's obviously something special about down, right?

So we conclude we cannot have a global Lorentz frame when we've got gravity. And in case you want to read a little bit more about it, I'll just give a highlight that this argument was originally developed by a gentleman named Alfred Shield. So that's unfortunate, OK, from sort of a philosophical perspective because the fact that we don't have a global Lorentz frame makes you think, oh, crap. What are we going to do with that framework we spent all this time developing? OK?

The existence of a Lorentz frame and many of the mathematics we've been developing for the past few lectures, they all center around things you can do in Lorentz frames. So it's kind of like, OK, so do we just throw all that out? Well, here is where Einstein's key physical insight came in and gave us the physical tools that then needed to be coupled to some mathematics in order to turn it into something that can be worked with but, nonetheless, really save the day.

So the key insight has there's a few steps in it. So first thing, I was very careful to use the word global when I ruled out the existence of Lorentz frames. But I'm allowed to have local Lorentz frames. Let's see. What does that actually mean?

So I'll remind that a Lorentz frame is a tool that we use to describe inertial observers. In fact, we often call them inertial frames because they're sort of the constant coordinates. They represent the coordinates of an inertial observer who happens to be at rest in that Lorentz coordinate system. So let's switch over. Let's start calling them inertial frames for just a moment.

So an inertial frame means that there is nothing accelerating, so no accelerations on observers or objects at rest in that frame. In other words, no forces are acting. Einstein's insight was to recognize that the next best thing is a freely falling frame. OK?

If you are in a freely falling frame, to you, it sure as hell seems like there's a force acting on it and there's some accelerations acting upon it. But let's imagine you're in this freely falling frame. And you have a bunch of small objects that you release near you. Let's see if I can find some bits of broken chalk.

So when you're in this really falling frame-- this was a lot easier when I was younger. All right, when you're in this really falling frame, OK, and these objects are all falling, there is no acceleration of these objects relative to you. The reason for this is that the gravitational force is proportional to the mass.

And so the fact that we have f equals ma. And fg proportional to mass means that all objects in that freely falling frame-- this is a term that I'm going to start using a lot. So I'm just going to abbreviate it FFF. If I get really lazy, I might call it an F cubed. Anyway, I objects that are in that freely falling frame, they're experiencing the same acceleration. And so they experience 0 relative acceleration at least in the absence of other forces, OK? Perhaps one of them is charging. There's an electric force.

Well, then it sort of suggests that the interesting force is the extra force provided to that relative to what we call the gravitational force that is driving this freefall. So I urge you to start getting comfortable with this idea because we're going to find this to be-- so one thing, which I'm going to do in a couple of minutes, is actually a little calculation that shows me I can, in fact, always find a Lorentz frame in the vicinity of any point in spacetime. OK?

And so we're going to actually regard that Lorentz frame as being the preferred coordinate system of a freely falling observer, one who is not accelerated relative to freefall at that point. We are not in freefall right now. The damn floor is pushing us out of the way. But in this way of doing things, we would actually regard us as being the complicated people, OK? Someone who is merrily plummeting to their death, they're actually the simple observers, who are doing what they should be doing, in some sense.

So the key thing, a way of saying this, so coming back to this and just one more point about that, because there are no relative accelerations, within this frame, objects maintain their relative velocities. That's basically the definition of something that is an inertial. If I have a frame where everything is-- two objects are moving with respect to each other at 1 meter per second, they're always moving at 1 meter per second, that's inertial. And so this sort of demands that we do all of our experiments in plummeting elevators or in space, right?

OK, obviously, you can't do that. So there's some complications we're going to have to learn the math to describe. But this is the way we're going to from now on think about things is that the freely falling frame is the most natural generalization of an inertial frame that we describe using Lorentz coordinates. Now before I discuss a few details about this, an important thing that is worth noting is that a very important aspect of any realistic source of gravity is that it is not uniform. OK?

Gravity here is a little bit stronger than gravity on the ceiling, OK? And it's a lot stronger than gravity in geostationary orbit. We call this variation in gravity tides. So tides break down the notion of uniform freely falling frames.

What this is telling us is that this idea, if I'm going to create this inertial frame that is essentially a freely falling frame, it will have to be a finite size. Hence, it's local and not global. OK, physically, this is telling me that, if I have a really tall freely falling elevator, and I'm here, I have one friend here and one friend down there, even if we start absolutely at rest, let's say we're at rest. We measure reference back to the walls of this elevator. It's a very stiff elevator. We will see the three of us diverge away from each other with time, OK?

Because this person, let's say earth is down here, this person is feeling the-- and using Newtonian intuition, gravity is a little stronger. Gravity is medium. Gravity is weak. And so as viewed in that freely falling frame, let's say we center it on me here in the middle, I will see the person in the top go up, the person at the bottom go down. Another way of saying this, so let me just say that, so a very tall elevator, we'll see a separation of freefall since gravity is not uniform.

But there's another way of saying this. And what this is telling me is imagine I made a spacetime diagram that shows trajectories of the three of us, so the three of us that are plummeting in this elevator, OK? What we find, if we make our trajectories as we fall-- in fact, let's just go ahead and make a little sketch of this.

OK, so let's say here is the middle. Here is the top. Here is the bottom. I'm going to draw this in coordinates that are fixed on the person in the middle. So the person the middle, in their freely falling frame, they think they're standing still. They see the person at the top moving up and away, the person at the bottom moving up and away.

It's a very similar story to the little parable I sketched of that light pole moving up. These are not congruent trajectories. Another way to say this is that, in any moment, if I think of these as world lines, I can, at any moment along the world line, I can draw its tangent vector. The tangents do not remain parallel.

Now I'm going to take you back to probably, I don't know when you guys all learned this, but when you first started learning about geometry. For me, it was in middle school. And when you first started learning about geometry, you were usually doing geometry on the plane. And they gave you various axioms about the way things behaved.

One of them is now known as Euclid's parallelism axiom. And it states that if I have two lines that start parallel, they remain parallel. If you dig into the history of mathematics, this axiom bugged the crap out of people for many, many years, OK? Because when you look at pictures on a piece of paper or on a chalkboard or something like that, it seems right, OK?

But it really can't be justified quite as rigorously as many other axioms that Euclid wrote down. And it turns out, there's an underlying assumption. The assumption is that you are drawing your lines on a flat manifold.

There's a ready counterexample. Go to the Earth's equator. You start, say, somewhere in Brazil. Your friend starts somewhere in Africa. Both of you stand on the equator and go due north. You are moving on parallel trajectories. You will intersect at the North Pole. OK?

Well, guess what? You ain't moving on a flat manifold. You're moving on the surface of a sphere. OK, curvature causes initially parallel trajectories to become non-parallel.

What this is telling us is that the manifold I'm going to want to use to describe events when I have gravity cannot be flat. I'm going to have to introduce curvature into it. Take us a little while to unpack precisely mathematically what that means. But this is sort of a sign that things got more complicated. And tides are the way in which that complication is being introduced.

So switch pages here. So what we have basically danced around and put together here is one formulation of what is known as the principle of equivalence. The principle of equivalence, or one formulation of it, tells me that over sufficiently small regions, the motion of freely falling particles due to gravity cannot be distinguished from uniform acceleration.

The physical manifestation of this is that if you're an observer who is in that freely falling frame moving due to gravity, you are also experiencing that uniform acceleration. So you see no net acceleration. This particular formulation of it is known as the weak equivalence principle.

I'm going to give you a couple variations on this in just a second. One thing that's important about it is, essentially, it's a statement that, if you think about Newton's laws, the idea that f equals ma and the force of gravity is proportional to m, it's a fairly precise statement that you think of that m as the gravitational charge. It's saying that the gravitational charge, the gravitational mass and the inertial mass are the same thing.

That's actually a testable statement. What you can do is look at materials that have very different compositions. In particular, what you want to do is look at things that maybe have different ratios of neutrons to protons or have highly bound nuclei with lots of gluons in them or various kinds of fields that you can imagine perhaps couple to gravity differently than the quarks do.

And so there are what are called free fall experiments to test this, which really basically just boils down to dropping lots of different elements of the periodic table and making sure they all fall at the same rate. And the answer is that they fall at exactly the same rate within experimental precision. And the last time I looked it up, well, what they do is they demonstrate the WEP, the Weak Equivalence Principle, is valid to about, I believe, it's about a part in 10 to the minus 13. Might be a little bit better now.

So bear with me just one second here. OK, so good. We've got this notion here. And you might now think, OK, great. I have a way of generalizing what the notion of an inertial frame is. Is this enough for me to move forward? Can we now do all of physics by just applying the laws of special relativity in freely falling frames or, putting it this way, our new notion of an inertial frame?

For two reasons that are closely related, this doesn't quite work. And it comes down to the fact that because of tides, we can't do that. What this basically means is that the transformation that puts you into a freely falling frame here is not the same as the transformation that puts you into a freely falling frame 10,000 miles over the Earth's surface, OK?

There are different transformations at different locations in spacetime. This comes down the idea that it's not a global Lorentz frame, OK? We had a sequence of local Lorentz frames that we have to link up.

So I'm going to do a calculation in just a moment where we explicitly show that we can always go into a frame that is locally Lorentz, but that there's a term that I'm going to very strongly argue is essentially the curvature associated with the spacetime metric that sets the size of what that region actually is. So what we are going to do, what is going to guide our physics as we move forward is, essentially, we're going to take advantage of a reformulation-- there we go-- of the equivalence principle, which I'm going to word as follows.

In sufficiently small regions of spacetime, we can find a representation or a coordinate system such that-- I don't want to cram this into the margin. So I'll switch boards-- the laws of physics reduce to those of special relativity. This is known as the Einstein equivalence principle. At least, that's the name that Carroll uses in his textbook. A couple books use different ones. But this is a very nice one because it really is the principle by which Einstein guided us to rewriting the laws of physics.

So we've got a weak equivalence principle, an Einstein equivalence principle. Just as an aside, there is something called a strong equivalence principle. We're not going to talk about this very much, but I'm going to give it to you because it's just kind of cool.

It tells me that gravity falls in a gravitational field in a way that is indistinguishable from mass. That's a little weird. I'm not going to explain it very much right now. A way to think about it will become clearer in some future problem sets.

Basically what it's telling us is that when you make very strong-- when you make any kind of a bound object, some of the mass of that object, in the sense of the mass that is measured by orbits, can be thought of as gravitational energy, OK? Energy and mass are equivalent to one another. So that gravitational energy should respond to gravity. And in fact, it falls just like any other mass, OK?

This is another one. This is actually very-- well, prior to September 2015, this was very difficult to test. There were some very precise measurements of the moon's orbit that were done to test this. And there were observations of binary pulsar systems that were done to test this. Now every time LIGO measures a pair of binary black holes, which are nothing but gravitational energy, and our theoretical models match the wave forms, sort of like, boom, equivalence principle. Drop mic. Leave room.

All right, so let's get precise. What I want to do is show that we can always find a local Lorentz frame. And what I mean by that is that I want to be able to show that I can always do a change of my coordinates, a change of my representation such that, at some point, I can convert the metric of spacetime into the metric we used in special relativity, at least over some finite region.

So let me define a couple of quantities and then let me formulate the way this calculation is going to work. Let us let coordinates with unbarred Greek indices be the coordinate system that we start in. And let's say that, in this representation, the metric is G alpha beta.

Let's demand that there exists a set of coordinates that I will represent with bars on the indices. When we transform to these coordinates, I want spacetime to be Lorentz at least in the vicinity of some point or some event. We'll call that event p.

So we will assume that there is some mapping between these coordinates, so there's no nasty singularities there. And it's nothing that we can't deal with. So we can do it in either direction, but let's say x alpha is x alpha of p written as some function of these guys. And so there's a transformation matrix between the two of them. It takes the usual form, OK?

So mathematically, the way I'm going to show this is I want to show that we can find a coordinate system such that, if I compute the spacetime metric and the barred representation, this is always going to be given by doing this coordinate transformation. Then I get the metric of flat spacetime over as large a region as possible. OK, we know we're not going to be able to do it everywhere.

What I want to show is that I can make that happen at the point p and that the functional behavior of this thing is sufficiently flat that it remains at this over some region. So let me just sketch what the logic of this calculation is going to be. So what we're going to do is expand.

So remember g is a function. All of these I's are functions. So what I'm going to do is I'm going to think of this whole thing as itself being some kind of a function. I'm going to expand in the Taylor series.

There's something really cool that we're going to have to do with this. I mean, if I do this in general, you're going to get something pretty vomitous It's a giant, giant mess. OK? You might worry, what the hell are we going to do with this?

OK, but there's something really cool that we can do. So we are given the metric g. But we are free to pick our coordinate transformation to be whatever we want it to be.

So what we are going to do is compare the degrees of freedom offered by the coordinate transformation. Which again, I emphasize, we are free to make that whatever we need it to be. So the coordinate transformation and its derivative, right? We're free to play around with this thing. It's under our control. To the constraints that are imposed by the metric and its derivatives, which we are given.

So what we're essentially going to do-- so in a moment or two, I'm going to write out very schematically what this coordinate transformation is going to look like when I do the Taylor expansion. And what I want to do is, essentially, just count up how many degrees of freedom the coordinate system offers me, count up the number of constraints that need to be matched in order to effect this transformation, see whether I've got enough. Assess what we learn from that, OK?

We just set up with a few more things to help define the logic. So I'm going to write G alpha beta as G alpha beta at the point p plus x gamma. And you know what? Let's do the expansion in the bar coordinate. Minus x gamma bar at point p times the derivative at point p. And then I'm going to get something that involves a complicated form of x squared and second derivatives. OK, this can keep going. But this is going to be enough for our purposes.

I am likewise going to expand my coordinate transformation matrix. Make sure I got all my bars in place. OK. OK. So let me go over to another board, just write down a few more important things. And then we'll start counting.

So first key thing, which I want to note before I dig into the calculation, is the metric at this point, its derivative, its second derivative. These have been handed to us. OK? So we have no freedom to play with these. These are going to give us constraints that we need to satisfy. The coordinate transformation in its various derivatives, we are free to specify these. These are our degrees of freedom.

OK, so let me just restate the calculation that we are going to do. Basically, we wanted to do this. Schematically, here's what this calculation is going to look like. And if you look at this right now, you might think this is going to be horrible. But we don't need to specify every step of this thing in absolutely gory detail.

What we want to do is make this, which I can write as this guy at point p, this guy at point p, this guy at point p plus a term that is going to be linear in the distance from the spacetime displacement away from point p. If you want to go and multiply this out, knock yourself out. But it suffices to know that you are going to get some new things that involve first derivatives of all these quantities. So this will involve these two things. And this will involve second derivatives of all these things.

Now what we're going to do is we're going to try to make what we would get multiplying all this out look as much like minus 1, 1, 1, 1 of the diagonal as possible. And like I said, it really just involves accounting argument. It's actually one of my favorite little calculations we do in this class because it's my favorite word for this is it really looks vomitous. But it's quite elegant.

And I have an eight-year-old kid. And I'm constantly trying to teach her, math is just all about counting. And this is a perfect example. This is really just about counting.

So let's look at this thing at 0th order. So what I'm trying to do is guarantee that I've got some l, which I can choose, which will make the components of the g at that point be minus 1, 1, 1, 1.

So my constraints, the conditions I must satisfy involve things that hit these 4 by 4 symmetric tensor components. So this guy is something that is symmetric. Nature has just handed me a symmetric tensor with a symmetric 4 by 4 object. Those are 10 constraints that my transformation must satisfy.

What I'm going to use to do this is this matrix at the point p. And I'll remind you, this matrix is just following set of partial derivatives at that point p. This is also 4 by 4. It is not symmetric. So this actually gives me 16 degrees of freedom to satisfy these constraints.

Ta da! We can easily make this thing look Lorentzian at the point p. In fact, we can do so and have six degrees of freedom left over. Any idea what those six degrees of freedom are? Remember you're going into a Lorentz frame. Yeah? Three rotations, three boosts. Exactly right.

Boom. So not only does it work mathematically, but there's a little bit of leftover stuff which hopefully hits our intuition as to how things should behave in special relativity. OK, that's not really good enough though, right? So I just showed that I can do it at this point. And if I move a centimeter away from that point, suppose, there's some really steep gradient, and then it goes completely to hell. OK? Then we're in trouble.

So we have to keep going. We have to look at the additional terms in this expansion. So when we do things at 0th order, we have now-- so g has been handed to us. We have now specified behavior of I at that point. So I has been completely soaked up. I don't have any more freedom to mess around with it.

When I go to the next order, OK, so the quantity that is setting my constraints is the first derivative of this. So these are four derivatives of my 10 metric functions. 4 by 4 symmetric by 4 components. So I have got 40 constraints I must satisfy.

OK, well, the tool that I have available, I am free to specify the derivatives at this point. Again, remember this is itself a partial derivative. So I've got 4 components alpha, 4 derivatives with gamma bar, 4 derivatives with mu bar. Partial derivatives, it should not matter what the order is. So this needs to be symmetric on mu and gamma.

So I, in fact, have 40 degrees of freedom at 1st order. Perfect match. OK? So I can make my coordinate system. Not only can I make it equal to the Lorentz form at-- I can make the spacetime metric Lorentz at the point p, I can also make it flat at that point p.

All right, now we're feeling cocky. So let's move on. How far can I go? I lost the page I need. OK, so now you have a flavor for what we're doing. OK? I want to count up the number of degrees of-- sorry-- the number of constraints that are in the metric object, the new metric object I work with at this order and compare it to the degrees of freedom in the coordinate transformation at this order.

So when I go to 2nd order, my new metric thing that I'm going to be messing around with is the second derivative. So I take two derivatives of g alpha beta at the point p. So this again needs to be symmetric in the derivative. So I have a symmetric 4 by 4 on these two guys. And the metric is itself 4 by 4, one alpha and beta. So there is 100 conditions, 100 constraints that we must satisfy.

So let's look at the second derivative of this. OK, so I'm going to move this down a little bit. This is now going to look like the third derivative when I do this transformation. So it's going to be the third derivative of x alpha with respect to mu bar, delta bar, and gamma bar.

So I've got 4 degrees of freedom for my alpha. And this must be perfectly symmetric under the any interchange of the indices mu bar, delta bar, and gamma bar. OK? This is a little exercise in combinatorics. So the number of equivalent ways of arranging this turns out to be n times m plus 1 times n plus 2 over 3 factorial. And I'm in 4 dimensions for n equals 4. Work that out, and you will find that you have 80 degrees of freedom.

So what we can do is we can always find a coordinate transformation that makes it flat. It makes it Lorentzian at point p. It is flat in that region. In other words, there is no first derivative there. Sorry. Flat's not really the right word. There is no slope right at that point.

But we cannot transfer away the second derivative. And so what this tells me is the coordinate freedom that we have means that we can always put our metric into the following form. I'm just going to write this schematically. Second derivative to the metric and sort of the quadratic separation in spacetime coordinate.

So a couple things about this are pretty interesting. So this is basically telling me that we can make it Lorentz only up to terms that look, essentially, like the second derivative of the metric. Well, the second derivative of any function tells you about the curvature of that function. So this word curvature is going to be coming up over and over and over again.

We are, in fact, because this is general relativity, we can't do everything simply. We're going to actually find that there is a rigorous way to define a notion of curvature that we're going to play with, which indeed looks like two derivatives of the spacetime metric. And we are going to find it has, I mean, it's actually going to end up looking like a 4 index tensor. OK?

So it's going to be an object that's got 4 indices on it. Each of those indices goes over the 4 spacetime coordinates. And so naively, it looks like it's got 4 to the 4th power independent components. 256.

It has certain symmetries we're going to see, though. And when you take into account those symmetries and you count up how many of those components are actually independent, any guesses what the number is going to turn out to be? 20. Yes, it exactly compensates for the number that cannot be zeroed out by this coordinate transformation. It's called the Riemann curvature tensor, and we will get to that fairly-- actually, really, just in a couple lectures.

The other thing which this is useful for is in all of my discussion of the equivalence principle up to now, there's been a weasel word that I have inserted into much of the physics. I always said over sufficiently small regions. OK? I say that trajectories begin to deviate from one another when they get sufficiently far away from one another. OK?

And you should be going, what the hell does sufficiently small, sufficiently far away, what do all these things mean? OK? And the issue is you need to have a scale. Well, this scale is going to be set by the second derivative of the metric.

So the size of the region over which spacetime is inertial in this coordinate system, in this representation is approximately so imagine it's 1 over-- we can think of d2g as being 1 over a length squared. So 1 over the square root of that gives me a rough idea of how long the curvature scale associated with your spacetime actually is. And it tells you how big your inertial region actually is.

So we're going to make this notion of curvature regress very, very soon. As a prelude to this, we now need to start thinking about how we do mathematics and physics on a curved manifold. So I'm going to start to set up some of the issues that we need to face.

So we're going to need to define what I mean by a manifold that is curved. So a curved manifold is going to simply be one in which initially parallel trajectories do not remain parallel. So an example is the surface of a sphere, as I illustrated. You start somewhere on the equator in Brazil. Your friend starts somewhere on the equator in Africa. The two of you start walking north. You are exactly parallel when you take that first step of the equator. But your trajectories cross at the North Pole.

Interestingly, an example that looks curved but is not, the surface of a cylinder. OK? If I take two lines, again, I need to imagine that this region continues all the way up here. I make two lines that are parallel to each other here and I have them extend around this thing, they would remain parallel the entire way out, OK? Another way of stating that is that you can always take a cylinder and, with an appropriate cut, but without tearing it, you can flatten it out and make it into a simple sheet, a perfectly flat sheet. You cannot do that with a sphere without tearing it in some places.

So what is going to begin to complicate things is that we want to work with vectors and tensors that live in this curved manifold. We haven't really thought too carefully yet about the space. And let's just focus on vectors for now. We haven't thought too carefully about the space in which the vectors actually live.

So implicit to everything we have talked about up until now is that we often regard vectors as objects that themselves reside in a tangent space. OK, if I'm working in a manifold that is flat, say it's the surface of this board, OK, and just two dimensions, every point on this board has the same tangent, all right? So if I draw a vector here and I draw another vector over here, it's really easy for me to compare them because they actually live in the same space that is tangent to this board.

If I'm on the surface of a sphere, points that are tangent to the sphere at the North Pole are very different from points that are tangent to the sphere on the equator. So it becomes difficult for me to actually compare fields when they are defined on a curved surface like this. So this makes it a little bit-- and so I'm just going to set up this problem, and I will sketch the issue. And then we will resolve it in our lecture on Tuesday.

This makes it a little bit complicated for me to take derivatives of things like vectors when I'm working in a curved manifold. So let's consider the following situation. So I'm going to define some curve, which I would just call gamma. It lives in a curved space of some sort.

I'm going to draw it here on the board. But imagine that this is on the surface of a sphere. OK? So here's the curve gamma. And let's say this point p here has coordinates x alpha. And this point q here has coordinates x alpha plus dx alpha.

Suppose there's some vector field that fills all this manifold. OK? And so let's say the vector a looks like this here at the point p. And it looks like this here at the point q.

How do I take the derivative of the vector field as I go from point p to point q? Well, your first guess should basically just be do what you always do when you first learn how to take a derivative. OK? So that would be a notion of a derivative. There's nothing mathematically wrong with that. OK?

But we're going to find that it gives us problems for the same reason that we ran into problems when we began working with vector spaces and curvilinear coordinates. OK? If I want this to be a tensorial object in the same way that we have been defining tensors all along here, I'm going to run into problems. So let me just actually demonstrate what happens if I try to do this.

So the key bit, the way we want to think about it is that the points p and q don't have the same tangent space, which is a fancy way of saying that, as I move from point p to point q, the basis vectors are moving. They're starting to point in different directions. So if this were to be-- like I said, mathematically, if you just want to get that derivative, that is a quantity which has a mathematical meaning, OK? But it's not the component of a tensor, which we have called out as having a particularly important meaning in this geometric construction of physics that we are doing.

And so if this were to be tensorial, then I should be able to switch to new coordinates, which I'll designate with primes, such that the following was true. The reason why this doesn't actually work is I'm going to demand that alpha-- excuse me-- that a is, in fact, actually already tensorial. OK? It's the compound of a vector, which is a particular kind of tensor.

So I'm going to demand that the following be true. And I'm going to demand that my derivatives, they are just derivatives. They do the usual rule Jacobian rule when I switch coordinate systems. And so skipping a line of algebra, which you can get in my notes, this is actually very similar to the notes I've already posted. When you work this out, you're going to find you get one term that's correct, but you get another term that involves a derivative of your coordinate transformation matrix.

And this is an actual term, spoils the tensorialness of this quantity. The way we are going to cure this is we are going to demand that if we want our derivatives to be derivatives that comport with a notion of taking tensor objects and getting other tensor objects out of them, before we take the derivative, we have to have some way of transporting the objects to the same location in the manifold, where I can then compare them.

So here's one example of a notion of transport that we could use. So here's a at point q. There's a at point p. Notion one that we'll talk about is known as parallel transport.

What that's essentially going to mean is I'm going to say, well, let's take one of these guys, either q or p. The way it's drawn in my notes, I've used q. But it could be either. And let's imagine I slide it parallel to itself until this gives me a alpha from q transported to p. And then I define a derivative with that transported notion of the vector.

Now notice I've called this notion one. You can take from that this idea of how I transport the vector from one place to the other to do this comparison. I cannot uniquely define it. There are, in fact, multiple ways you can do this.

We're going to talk about two that are useful at the level of 8.962. In principle, I imagine you could probably come up with a whole butt load of these things. These are two that the physics picks out as being particularly useful for us for the analysis that we are going to do. OK?

So we'll pick it up there on Tuesday. We're going to start by coming back to this notion of I want to differentiate a vector field in a curved manifold. And let me just state before we conclude, now that I've transported a from q to p, they share the same tangent space. Since they share the same tangent space, I can compare them more easily. That allows me to make a sensible derivative that respects all the notions of what a tensor should be. We'll pick it up from there on Tuesday.