

[SQUEAKING]

[RUSTLING]

[CLICKING]

SCOTT

OK. So let's just dive in. And I'll just quickly recap where we were last time. So what I

HUGHES:

did last time was, we are now beginning the adventure, so to speak, of looking at how one examines-- how one finds and then studies the properties of solutions to the Einstein field equations. In other words, given a particular source, what is the space-time that emerges? What is the space-time that is consistent with that source of gravity in Einstein's relativistic theory here?

So we begin-- we're going to look at this in a couple of different ways. And the first one is where we're going to-- what we do-- linearize the Einstein equation. So we're going to imagine that space-time is a small deviation from the flat space-time that contains no gravity. And we are going to expand the Einstein field equations, discarding all terms that are non-linear in the perturbation away from flat space-time.

So last time we did this, we developed our field equations, we recast them in a form that's particularly convenient for making solutions, and we noted that the coordinate degrees of freedom can be thought of as a kind of gauge freedom that, in some ways, it's very useful for us, because it allows us to write the equations in a form that is particularly amenable to solving, but can also be a little bit dangerous, because you can sometimes wind up with a solution whose physical character is masked by the nature of that gauge. I went through a somewhat advanced topic for the purposes of 8.962, and I do not expect students to follow this in detail. But it's good to at least survey this and understand what was going on with this.

What I did was I decomposed the perturbation to flat space-time into irreducible degrees of freedom. I sort of chose a time coordinate. I then imagined the other three coordinates in our space, and I looked at how, having chosen time, the different degrees of freedom decompose in two scalars with respect to coordinate rotations, vectors with respect to coordinate rotations, and two-index tensors with

respect to coordinate rotations. I then took those, particularly the vectors and the tensors, and I further broke them up into degrees of freedom that are either divergence-free vector fields or the gradients of scalar fields.

And in doing this, what we found was that the 10 degrees of freedom that are specified-- excuse me, the 10 degrees of freedom that are present in the perturbation the space time, can be broken down into six functions that are gauge-invariant plus four gauge degrees of freedom. Applying a similar decomposition to the source, we found that those six gauge invariant degrees of freedom are governed by four Poisson type equations. And I will refer you to the previous lecture for exact definitions of these things, but there is a scalar field ϕ that is related to the mean density of energy in your space-time and pressures in your space-time-- sort of isotropic stresses.

This DTS is an additional degree of freedom that is defined in that previous lecture. You find that a Poisson equation for a scalar field I called θ is directly related to the energy density. There are two degrees of freedom associated with sort of off diagonal terms in the metric-- the time-space terms in the metric. And they are related to the momentum flow. These are only two degrees of freedom, because the ψ is a divergence-free field.

And finally, what we found is if you take the spatial piece of the space-time metric, you project out the bits that are transverse-- in other words, it's divergence-free. We will come to a more physical meaning of this a little bit later in today's lecture. And you require that this be traceless-- then those components of the spatial space-time perturbation are already gauge invariant, and they're related to sort of anisotropic stresses in our source.

They are governed by a wave equation. And as such, these represent a radiative degree of freedom. Because this is [coughs] (excuse me!)-- because this is trace-free and divergence-free, naively you would think that this has six numbers in it. A symmetric 3 by 3 matrix is how you would represent it. But because it is trace-free and because it is divergenceless, there are four constraints. And so those six independent numbers are reduced to two. And as we'll see in this lecture, those two degrees of freedom correspond to the polarizations of gravitational waves.

Today's lecture is all about understanding these guys. Given our kind of funny schedule that we're on today, I got up early and have not had time to take a shower, so I am wearing my LIGO hat. And what we're going to be doing today is the foundations of how one calculates this and why one builds an observatory like LIGO to measure these things. All right. So let me just write that out.

So our goal for today-- I should say our goal for this lecture, because this is the first of three lectures I'm going to record today. Our goal is to understand h_{ij} TT in terms of quantities that we can observe and also to understand how to compute this field-- this h_{ij} TT given us source. So let's begin by just assuming. Let's begin by focusing on how we can understand what h_{ij} TT actually means-- how we can go and measure this. So what I'm going to do is just we'll start out by let's just hand ourselves an h_{ij} TT and study its properties.

So what I'm going to do is begin with a space-time perturbation that has everything except h_{ij} TT equal to 0. So this is something that has ϕ equals 0, θ equals 0, ψ equal to 0. I am going to pick h_{ij} TT. If we have a form such that it only depends-- it's a function that depends only on the combination T minus z . So this represents something propagating-- as we see, it actually describes radiation. in the z -direction. So this assumed form can be represented by the following matrix.

So if you take a look at this, I'm going to justify a few details of this later. There's a few constraints. So it looks like I have four numbers here. Because my space-time metric is always symmetric under-- we'll just say symmetry-- under exchange of indices, h_{xy} is the same thing as h_{yx} . So these two things are the same.

And from the requirement that this construction be trace-free-- so from the way I've written this here, that tells me that h_{xx} plus h_{yy} must equal 0. So these two quantities on the diagonal are not independent. They must be of opposite sign. So there are my two degrees of freedom. Why have I set the z part equal to 0? We'll look at that in a little bit more detail a little later in the lecture.

But essentially, it comes from the fact that I'm making it propagate in the z -direction. What we're going to see is that if-- whatever direction the wave is propagating in, my field has to be orthogonal to it. And we'll prove that a little bit more carefully later.

So, for now let's just take this as handed to us. We can see this is at least consistent with the properties I've written down here, from the fact that it depends on T minus z , it's not hard to show that this thing will satisfy the wave equation. So we have a good source for exploration in front of us-- excuse me-- a good space-time for us to explore the properties of. So let's begin by saying great. One of the most important things we've learned how to do is, given a space-time, I can compute how a freely falling body moves in that space-time.

So let's consider what happens if I have freely falling bodies in a space-time, in a space-time such that this-- in a space-time that is this-- so I am not saying this very fluently. So this is what we want to do. We're going to consider freely falling bodies in a space-time that involves this $h_{\mu\nu}$.

And the way I'm going to do this, my first pass is to sort of say, OK, I know all about the motion of freely falling bodies in space-time. They are geodesics. Geodesics describe such bodies. So let's go ahead and look at the geodesic equation for this setup. Now I wanted to-- let's give ourselves a couple of conditions to try and understand this.

Let's imagine that these are-- we're initially in a space-time that is just flat space-time. So we're going to imagine that $h_{\mu\nu}$ is 0 initially. And perhaps it then-- we're sort of imagining a space-time in which there's just you falling freely in empty space. Absolutely nothing going on here. And this $h_{\mu\nu}$ field passes over you. A wave packet passes over your freely falling particles.

So $h_{\mu\nu}$, not equal 0, comes along, and passes over your particles. I'm going to imagine that these bodies are initially at rest. So when I'm in this regime where $h_{\mu\nu}$ is equal to 0, these guys are just sitting stationary with respect to whoever is making these measurements. So we're going to say great. I've got these bodies sitting at rest in what is initially empty space-time. A wave passes along. What happens to these bodies?

So if these guys are at rest, then I know that, at initial times-- so let's just say that this comes along, and it becomes $h_{\mu\nu}$ not equal to 0 at some time-- what is defined as the origin of time. So t equals 0 is when the non-zero wave comes along. So my space-time is just $\eta_{\mu\nu}$ for all earlier times. My particles are sitting here

at rest at all earlier times, then $h_{\mu\nu}$ comes along.

So what we want to do, then, is look at how does this guy behave at $t = 0$. Now, initially, the only component that will get picked out of this is the 00 component. And, working to linear order, I can write this. Just wanted to write one more step in here. Let's say that this, I'm going to write as $\eta_{\alpha\beta}$. Go back to your earlier lecture, when we defined the Christoffel symbol to look up the definition of this guy.

So these are all the only components of the space-time that enter into this Christoffel symbol. Every single one-- let's go back and take a look at the space-time that we're looking down here. Notice this involves $\beta_{0,0}$, 0β , and 00 . That is the space-time parts, all of which are 0. The other space-time parts, all of which are 0. And the time-time part, which is 0.

So the initial acceleration is 0. That means $du^\alpha/d\tau$ is equal to 0. Which means that my particle begins on a geodesic, and it stays on the geodesic. In these coordinates, it does not move at all.

This seems to suggest that gravitational waves have no effect. I have a small body here. Gravitational waves come along. It's unaccelerated. So what's this for? Well, if you think about this for a second, what you should convince yourself is that the calculation I just did, it's probably the first one you would think of doing when you're trying to ascertain how does a body move in a particular space-time. But, in some ways, it was kind of a dumb calculation to do.

So let's bear in mind, what this is doing is tell me the geodesic equation describes motion with respect to some given coordinate system. Satisfying this equation is just saying, yup, you're following a geodesic. And guess what? These are-- this radiation is a form of gravity. Even if the body responds to it, it is still going to follow a geodesic. It's going to be in free fall. The gravitational field may be a little bit different from the ones that you're more used to from your previous experiences with gravity, but it is still just following the free fall trajectory that gravity demands.

So free fall in these coordinates means the body is remaining fixed in these coordinates. This is one of the most important lessons in general relativity. When you ask a statement that depends on the coordinate system you've written down,

you have to understand your coordinates. And I would submit we don't really understand the coordinates that we wrote down here with this. These coordinates essentially follows the body. Even if the body is moving, they essentially wiggle right along with the body itself.

So if you want to understand what the impact of gravitational waves is on things, you have to try to formulate what you're able to measure using coordinate-independent language as much as possible. At the very least, you need to understand your coordinate system a lot better than I did in setting up this calculation over here.

So let's reframe the way we do this analysis. Let's imagine-- so one of the things that we learned in thinking about how gravity works in general relativity is I can always, if I'm just focusing on a single body, a single point mass, I can always go into freely falling coordinates. And space-time is essentially the space-time of special relativity. I can find a representation so that everything looks just like special relativity near there. Gravity is basically completely thrown away, at least as far as the mathematical representation is concerned. I'm in free fall. I don't really see much.

If I wanted to see the impact of these waves, what I need to do is think about how things behave over some finite region. I wanted to put together a setup such that I have two bodies, and the relative effect of things between those two bodies will allow me to see the effect of space-time curvature in action.

So let's now-- with this in mind, let's consider two nearby bodies that each follow geodesics. So body A is right here. And I'm going to define body A as living at the origin of my coordinates. And we've got a body B over here. And we're going to say that this guy is displaced from the origin by a small amount, some epsilon, in what I will call the x direction.

And what I want to do is say, OK, I've got these two guys. They are moving through a space-time. Let's make this be the space-time that we had just discussed a moment ago. So they are each moving on geodesics of this space time. We know that if I set up my coordinate system in the way that I've done it here, I am not going to see any acceleration on each of these bodies. Let's now try to focus on some kind of a quantity that tells me about the separation between these two. I'm going to do this

in two ways.

So the first thing which I'm going to do is say, let's imagine that I am bouncing a light pulse back and forth between A and B. What I'm going to do is write the 4-momentum associated with that light. So recall I can't really define a 4-velocity for light because proper time is not defined. But I can introduce some kind of an affine parameter such that d by $d\lambda$ tells me about motion on the world line that this light follows. And this is just bouncing back and forth between these two. So it moves in the x direction.

Because this is light, $p \cdot T$ equals 0. So these components, the 4-momentum associated with the light, the $dt/d\lambda$ and the $dx/d\lambda$, they are related to each other like so. And look. The space-time metric, the gauge-invariant piece of the space-time, $h_{\mu\nu}$, it is appearing in here.

So let's say, OK, what I am interested now is computing what-- suppose I'm sitting at A, and I want to compute the time, according to A's clock, that it takes for a pulse of light to travel from A to B. Maybe I'll make it bounce back. Maybe I want to know how much time it takes for the light pulse to go from A to B and back to A.

So let's turn this into an equation for the interval dt . So dt , I can do a little bit of this trick called division. Here is-- question is, did I do it correctly? Yes. So I'm doing the old physicist's trick. I'm just imagining that an interval of time-- an interval of affine parameter $d\lambda$, the amount of dt that accumulates as it moves to an interval dx . They are related by so.

And, remember, I am assuming that the perturbation from space-time, from flat space-time, is small enough that I can always linearize. So this can be-- I can take that square root using the binomial expansion. So suppose I want to integrate up the time it takes for the thing to go from A to B, and then back from B to A.

Well, this is going to be given by integrating this guy's motion from 0 to ϵ , like so. And then integrating it back again. My dt is in here. It's going to switch direction when I do that, so I'm going to flip the sign on that dx .

I'm leaving it written in this form, so you have to be a little bit careful. You guys will play with this a little bit more on a problem set. Bear in mind that as I move through

x, the time argument of these guys is changing. And so that's where there's a little bit of work that you'll explore in a problem set. But I'm writing it in this form so that we get some-- naively, you might think, ah, I'll just flip the sign and I'll combine these. In a certain limit, that is indeed a good approximation. You got to be a little bit more careful in general though.

This is the time it would take for light to go from A to B and back in the absence of this gravitational wave, this hx_{tt} . Notice there is an inference on the time of arrival of these pulses that depends on hx_{tt} . Time of arrival of the pulses depends on this particular piece of the space-time metric.

One could imagine, if you had a precise clock, and you sat there and you timed the arrival of these things-- let's say you sent a pulse of light out-- bloop-- and you kept track of how long it took with a very precise clock, you could look for-- in the absence of a gravitational wave, that would be, if you imagine that your pulse was sent out perfectly periodically-- let's say every millisecond or something like that, you sent out a little pulse-- if it arrives back and your arrival pulses were just as spaced by 1 millisecond, you would say, great, I am sitting in empty space-time.

But if you found that their arrivals varied, you would think, hmm, there's something else going on here. Gravitational waves is one of things that could lead to that variation. So this, at least, is kind of a proof-- this calculation I just did-- you guys are actually, on one of the upcoming problems sets-- I believe it's problem-- I can't remember if it's 7 or 8-- you do a little analysis where you will use this to understand how a detector like LIGO actually functions to measure out the effects due to a gravitational wave on the arms of a detector like this.

Let me just do one more way of understanding how it is the gravitational wave leaves an imprint on these two separated particles. So when you look at that first naive calculation I did, the fact that I got no effect, I sort of joked that we did a dumb calculation. To be fair, it wasn't that dumb.

What we're-- the first thing you usually think of doing when you want to understand the motion of a body in space-time, is you look at geodesics. Geodesics just tell me that this guy's acceleration with respect to the free fall frame is 0. And that's exactly what we ended up getting. It's basically just saying, ah, even in a gravitational wave,

my two bodies are moving on geodesics.

If I have two separate bodies, like I have in this calculation where I just calculated the variation in the time of arrival of light pulses, one of the tools we've learned is that I can have two geodesics, two free fall trajectories that deviate from one another. So let's look at the geodesic deviation of the free fall world lines of these two bodies. So I'm going to-- as I did when I did my geodesic calculation, I'm going to take them to initially be close enough that they have the same 4-velocity.

And if you want to be careful about counting orders here, you can imagine, when the gravitation wave comes along, that one or both of them might pick up a correction of this that's of order h . Depends on what kind of coordinate system you pick up. But the key thing is that they start out like this. The gravitational waves may come along and it will move them around a little bit. But the amount it's going to move them, the amount of velocity that they're going to pick up, and the displacement they're going to get from their initial position if they're at rest, it can only be of order of the h that you're trying to measure there.

So I'll remind you, the equation of deviation, equation of geodesic deviation, it defines a sort of proper acceleration acting on-- whoops, that's a typo. It is the proper acceleration on a vector ξ that defines, in a precise geometric way, the separation of two geodesics. And what it looks like is the Riemann curvature of your space-time contracted on the 4-velocity of your bodies, and that separation vector ξ .

So, linearizing everything, bearing in mind that R will itself be of order h , and that any difference between the proper time along these things and the coordinate time of a particular clock is also going to be of order h , they are small enough that this equation becomes-- Something that looks like this. Notice I'm only looking at the spatial components of the separation vector ξ . And that's because, since I'm working, since this is what my 4-velocity looks like, I'm going to pick out the 00 components of these things.

And, don't forget, Riemann is antisymmetric when I exchange either indices 1 and 2 or indices 3 and 4. And because of that, the time-like component of this guy is going to be an uninteresting thing. I will just get that is equal to the acceleration of the

time-like component is 0. So it doesn't really matter.

So this is telling me that the geodesic separation of these two world lines is-- it's a second-order equation in time, and it's proportional to the separation itself. We need to work out all of these components. That is a fairly straightforward exercise, and I will just report to you the results.

A brief aside. If you want to sort of verify the results I'm about to write out here, you could go into-- you could look up the definition of the Riemann tensor. Bear in mind that you are doing things at linear orders. You can throw away the terms of, like, connection times connection. It's just derivatives of the connection that are going to matter at linear order. But, even so, it's a bit of a grotty calculation.

I am-- as I begin to reorganize the problem sets associated with this class under the new system that MIT is currently operating on, I will soon be releasing a few Mathematica notebooks that are very useful tools for doing tedious calculations, like computation of Riemann curvature tensor components. So I will put a few of those things up there. If you would like to validate some of the things that I am computing, that is where you will find a good way to test this.

So, anyway, the non-zero Riemann components. So at linear order, you have R_{x0x0} . Bear in mind that, linear order, you are raising and lowering indices with the metric of flat space-time. So that's the same thing as this. And this turns out to be minus 1/2 two time derivatives of h_{xx} . $Y_0 Y_0$ -- that's your yo-yo component-- pretty much looks exactly the same, but it's two time derivatives of the yy piece.

Remember though-- I've erased it a while ago-- there are some constraints that we found due to the fact that my original $h_{\mu\nu}$ has to be-- its spatial piece has to be traceless. In particular, we found h_{yy} was minus h_{xx} . So I can remove that additional function, and just leave things in terms of h_{xx} . And our final non-zero component looks like this. All others are either 0 or related by a symmetry.

So notice, although I've written down three different complements here, there's really only two, because the symmetry that's-- it's not a symmetry of Riemann, but the symmetry associated with the trace-free nature of h_{ijtt} made these two equal to each other. So this and this are really the only independent degrees of freedom here.

All right. Let's plug this into the equation of geodesic deviation, this guy, and see what we get. And what we get is an equation governing the acceleration of the x component of deviation. Here it is. Let's make a couple of assumptions and try to understand what this is telling us.

So what I'm going to do is-- let's assume this is-- the first assumption is necessitated by the idea that I can do linearized gravity. So I'm going to assume that these guys are all very small. And I will imagine that my displacement vectors have some piece which-- basically, it's time constant. It does not change. And that there's a bit that will be of order h that describes how the separation of my two bodies responds to this incoming field.

So what I'm going to do, when you go and you plug this in, assume that this guy is constant in time, and this one varies. What I'm going to do is imagine I don't just have two bodies. Imagine I have a ring of bodies. So suppose I have got my coordinate axes here. Horizontal is x , vertical on the board is y .

And so imagine I have a ring of bodies, like so. And let's consider the following limit. Let's imagine that my h_{xx} is some function that is sinusoidally varying in time, and my h_{xy} is 0. I just want to call that the-- as we saw, there's really only two degrees of freedom in this field. And in the way that we formulated it, it's going to be encapsulated by these two functions, xx and xy . So let's just look at one. Let's just look at the influence of the h_{xx} .

So when you do this, you essentially have the second derivative of $x_i x$ looks like two derivatives of your x field, your xx field times $x_i x$. And two derivatives, your sort of acceleration in the y direction, is minus those two derivatives of your field times $x_i y$.

So what this is going to do, if you think about the way this is going to act on these things, you're going to get a sinusoidal acceleration associated with the displacement of this little sea of freely falling bodies as the gravitational wave comes along. And it's going to do so in such a way that the displacement in the x direction does the opposite of what the displacement in the y direction is doing.

So as time passes, you'll find that-- so, initially, let's say it stretches along x , but then squeezes along y . Then it's a sinusoid, so, a quarter of a cycle later, it's back to

being a circle. And, a quarter of a cycle later, it stretches along y and squeezes along x .

Let's consider the limit now where h_{xx} is equal to 0, and h_{xy} is some sinusoid. So I start out with my circle. Look at your equations over here. Think about what they're going to do. Now we're going to find that the change in x and the change in y is correlated. So this guy stretches it out into an ellipse along a 45-degree line. A quarter of a cycle later, it's back to being a circle. And a quarter of a cycle later, it does the same thing but with opposite sign.

These two fields, my h_{xx} and my h_{xy} , are what we call the fundamental polarizations of a gravitational wave. Polarization states are only defined with respect to a particular set of basis vectors. So if I make my basis vectors be this x and y , we call this one, for reasons that I hope are obvious, the plus polarization.

The gravitational wave acts as a tidal stretch and squeeze. Notice what's going on here. I stretch along this axis, squeeze along this one. A quarter of a cycle later, I squeeze along this one, stretch along that one. So I get a tidal stretch and squeeze, where the main directions along which the stretch and squeeze are happening is plus shaped with respect to these axes. This one is known as the cross polarization. You also have a tidal stretch and squeeze, but the lines of force in this case are along axes that have an X shape.

So at the level of 8.962, this is all we're going to say about the way that these things act on-- how gravitational waves act and what their observables are. And I deliberately chose to wear my LIGO today, because what LIGO does is look for these tidal stretches and squeezes. And the way it does so-- so, obviously, it's technologically not really feasible to set up a ring like this. But what we can do is sample the ring here and here. Actually, it's a little bit better to do on the plus polarization plot.

So you make an experimental setup that samples the ring here and here. And if a gravitational wave comes along and it's lined up right, you see, ah, look at that. I'm getting that stretch and squeeze. The way you actually measure it out is essentially a timing experiment, exactly like this little calculation I did here. You bounce light between mirrors, and you very carefully time the round-trip travel time from one

mirror to the other.

And what's particularly cool is the experimental aspect of this. So how does one measure the time so precisely? Well, you use a laser beam as your clock. And what's beautiful about doing that is the way that you can actually sort of do the metrology at the level necessary to get the precision to measure your h_{xx} 's in here is by interference. So you can treat the incoming laser beam as defining your time standard, and then by interfering the light that has traveled down the arm and bounced back with that laser beam, basically the laws of nature allow you to automatically check and see whether there is a shift in the time of arrival associated with the action of the gravitational wave.

So I imagine there will be some questions about this. There is a homework exercise that allows you to sort of test drive some of these concepts and allows you to see how it is that LIGO works. It really just sort of touches the-- it just scratches the surface of what an amazing experimental feat it is. But it should give you a good idea as to what goes into that.

Now, there is an important part to all this that we've not talked about yet. I handed you h_{ij} . Really, I handed you $h_{\mu\nu}$. The big thing that we care about next is, given a particular source, how do I compute this radiation field? So bearing in mind that h_{ij} , this is the gauge-invariant radiative degree of freedom in my space-time, how do I compute this guy?

And what I'm going to do is borrow sort of one of the most important lessons of that gauge-invariant formalism. So we need to be cautious of the fact that, in some gauges-- for example the Lorentz gauge, which is so convenient for solving the Einstein equation-- components of the metric appear to be radiative whether they are or they aren't. This is a consequence of our gauge.

But we also know only h_{ij} is radiation in all gauges. So that was what the previous lecture was about. As I've emphasized, it was a slightly advanced one. I don't expect everyone to follow all details. But if you can sort of grok these two main concepts, you're ready to now exploit this in the following way.

What you do is, step a, compute-- find some metric perturbation in any convenient gauge by any convenient method. Once you have that, extract h_{ij} . You're done. So

that's what we're going to do right now in the last 20 or so minutes of this recorded lecture.

So the way we're going to do it is let's take the linearized Einstein equation in Lorentz gauge. As we showed in-- I believe it was two lectures ago, the Einstein field equations in this case reduce to a flat space-time wave operator on the trace-reversed metric perturbation, being up to a factor of minus $16\pi G$ of the stress-energy tensor. And this has an exact solution.

So I'll remind you of this notation. So this is the field at time t and spatial location x . x and t -- well, let's write it like this. t and x are my field point. That is where I measure my field. x' is my source point. That is where I am looking at my particular contribution to the integral. As we're doing this calculation, there was a t' , but the radiative Green's function turned that t' into this retarded time. So what we are seeing is that the field at t depends upon what's going on at the source at t minus the interval of time it takes for information to propagate from x' to x .

So what I'm going to do is, first, I'm going to introduce a couple of approximations that reflect the reality of the conditions on which we typically evaluate this integral. And then we're going to talk about how to solve this thing, and, given that we have now solved it, how to project out the gauge-invariant degrees of freedom associated with the radiation. Whoops, pardon me. Let's put this up.

So first is I introduce this approximation. It's worth bearing in mind that when you are calculating this, x minus x' , which is kind of like the distance from a particular point in the source to where you're making your measurement, it is generally a lot larger than the size of the source. So the kind of situation we're thinking about is, here we are, making our measurements at some point x here. Here's my source, and here is a point x' inside the source.

So the thing to bear in mind is that, for the kind of calculations where that-- when one is actually doing this, x minus x' -- well, if I think about the kind of measurements that LIGO is doing, for example, that is typically hundreds of millions to billions of light years. The size of the source itself is of order tens to hundreds of kilometers. Bit of a separation of length scales there, you would say.

So what I'm going to do is approximate the solution of that integral, and, for pedantic sake, I'm going to use an approximately equal. What I'm going to do is replace x minus x prime with r . Let's just say that this distance here is effectively the same for all points in the source. I can pull that out of the integral.

If you like, you can refine this. There's a more careful expansion. And if you've done any electrodynamics, you should be familiar with this. And it uses the fact that 1 over x minus x prime can be expanded in Legendre polynomials, in a sum of Legendre polynomials. So if you do that, you can introduce refinements to what I'm about to derive right here.

What this does, then, is it defines a multipolar expansion. I am not going to do that. I'm just going to look at the very first term. And you can sort of think of that as saying that I am just doing the leading, the most important multipole. Additional multipoles beyond that are certainly important, and people who work in the field of gravitational radiation, we certainly do not neglect them. In fact, much of my current research is based on sort of high-order calculations associated with looking at behavior of some of those multipoles. This will be fine for initial pedagogical purposes.

Next, we're going to use the fact that only the spatial-- whoops. That's not supposed to be an h tilde. My apologies. That's supposed to be h bar, because this is the trace-reverse thing. Weird times, folks. Where was I at? Yes. So I am only going to care about the spatial pieces of this. You might protest, hey, isn't it the spatial part of the metric, not the trace-reverse metric, that matters? And you would be correct.

If you do this carefully, to be perfectly blunt, you have to sort of finish the calculation before you can kind of justify the step I'm about to make. And we're not going to have time to do that very carefully, but I'll make a comment about it. For now, it suffices to say we only care about the spatial components.

So I'm going to take my α β over to an ij . So I now have something that involves the integral of the spatial pieces of the stress-energy tensor over my source. And now I'm going to take advantage of an Easter egg that was hidden in an early problem set. Back on problem set 2, I asked you to prove a result that I called the tensor virial theorem. One of the reasons why I assigned that was that I knew I

was going to need it in this lecture.

And what the tensor virial theorem told us is that if I integrate that, this ends up being equivalent to $1/2$ two time derivatives. It's equal to the integral of two time derivatives-- excuse me. It's two time derivatives of the integral of the 00 piece of my stress-energy tensor. So it's the second moment of that, x_i prime, x_j prime.

This is great, because what I've got in the left-hand side here, that's exactly what my integral is. Actually, you know what? This is going to give me a result that's sufficiently clean. I want a new board.

So when that is done, what I finally see is that my h_{ij} , trace reversed, is $2G$ over r two time derivatives of a tensor that I will call l_{ij} . This is what I get when I integrate T_{00} , two moments of T_{00} . This is called the quadrupole moment of the source. The asterisk is because, as you guys will see on one of the problem sets, there's a slight tweak that goes in the exact definition of the quadrupole moment which has to do with the behavior of the trace of l_{ij} . I will just leave this like so for now, and leave you guys to explore that on a problem set.

So we're almost done. What I now need to do is pull out the transverse and traceless piece of this. Let's consider the transverse condition first. So, to do the transverse condition, I'm going to take advantage of the fact that, far from the source, \bar{h}_{ij} is equal to 0, because I'm in a region where the source is equal to 0 when I'm far away from it. This suggests that what we ought to do is expand our solution plane waves. So what I can do is write \bar{h}_{ij} . And I'm going to choose my indices carefully here, call it j_l . It's some amplitude that does not depend on space or time, $e^{i(\omega t - \mathbf{k} \cdot \mathbf{x})}$.

So this amplitude may depend on ω , may depend on k . It does not depend on t or on x . And the reason I changed my indices from i and j to j and l , and I did not use the index k , is just to avoid confusion between square root of minus 1 and the wave vector. To be absolutely careful, this is a single mode. If you like, maybe you should think about this as something that I have-- I'm doing a sum or an integral over ω and k in order to reconstruct the whole thing. This is just a single Fourier mode in the wave field.

So if I want this to be transverse, that condition is that the divergence of this guy be

equal to 0. But when I take the divergence of this, this becomes-- I pull down a factor of minus i , and I pull down the j -th component of my wave vector k . And I get my wave-- my Fourier term.

That can be thought of as-- so we can clear out the factor of minus i . The transverse condition boils down to saying that k_j -- whoops. No, that's right. That's right. k_j contracted on the wave field is equal to 0. In other words, the wave field is orthogonal to the wave vector.

If you have studied electrodynamics, this should be familiar to you. This is the same-- it's exactly the same as the idea that when I have an electromagnetic wave, my E field and my B field are orthogonal to the direction of propagation. Let's define a unit vector n_j to be k_j but just normalized. So my wave is propagating in the direction of the unit vector n .

Go back to another Easter egg that was hidden on an earlier problem set. You guys studied projection tensors. If I introduce the projection tensor P_{ij} , P_{ij} will project any spatial vector. It will project out the components normal, orthogonal to the direction of propagation n . You can in fact think of P_{ij} as defining the metric for the subspace that is orthogonal to n .

So imagine that your wave is propagating in the n direction. There's a plane being carried along with it. The plane is orthogonal to n , and P_{ij} is the metric you use to define spatial distances in that plane. Using P_{ij} , we can now make my h trace-- excuse me. I can make it transverse.

All I need to do-- I'll put one T on this to denote Transverse. Whoops. All I need to do is project each of its indices within that tensor. So we're almost there. It's transverse, but not yet trace free. So if I take the trace on this guy, taking the trace on this guy amounts to-- I am going to require that this be equal to 0.

Now, suppose I have some tensor that is not trace free, and I want to remove its trace. I'm going to make this a two-index tensor. So I'm not going to prove the following result. You can easily verify for yourself. The formula for doing so is that a_{ij} trace free is $a_{ij} - \frac{1}{N} g_{kl} g^{kl} a_{ij}$. G is the metric that describes the manifold in which this tensor lives. N is its dimension.

Now, for us, the tensor field whose trace I want to remove is h_{ij} . So \bar{h}_{ij} is the field whose trace I wish to remove. This lives in the two-dimensional manifold whose metric is the projection tensor P_{ij} . And we are working in a space where raising and lowering indices, there's really no important distinction between the two of them, because we're working in linearized theory.

So I remove the trace by taking this guy and subtracting off-- pardon me just one moment. Yes, sorry. Subtracting off like so. Pardon me. It should have a T on it. Plugging in my various definitions, this turns into $h_{lk} P_{li} P_{kj}$ minus $1/2$. So you might want to go and check some algebra there, but this is all-- I'm very confident in this. There's an overall factor of \bar{h} .

And the last thing which I will note is that, once I have removed the trace, this bar becomes meaningless. It's trace free, so the trace reversing of it has no action. Now, you should be a little bit careful about that. There's a few different notions of trace that went into this definition. Without proof, and just stating that by being a little bit more careful of some of the definitions we put down, the difference between-- when you count this up, you'll find that you may be making a small error, but it is on the order of terms that are quadratic in the perturbation.

Putting all these pieces together, we finally obtain our transverse and traceless metric perturbation. Which looks like this. This is a result known as the quadrupole formula for gravitational waves. This is-- so given-- just take a stock of the final result.

Suppose you have some dynamic gravitating source. This is telling me, what I do is I compute the quadrupole moment associated with that source. I take two time derivatives. I hit it with this combination of projection tensors, which has to do-- that tells me things about the geometry of the source. That's saying that my source, the waves that come from the source propagate along a particular direction to me. This picks out the components that are transverse with respect to that direction and are traceless, so that I don't have any gauge degrees of freedom in what results.

Multiply by twice Newton's constant, divide by distance. And this is a wave field that our friends over in Building NW22 can now measure pretty much every day. You're going to do a few homework exercises associated with. One thing which is worth

commenting on before I move on is, when you keep factors of c in this, that G goes over to G over c to the-- wait a minute. Let me double-check that. Is it G over c to the fourth or G over c squared?

Anyhow, it's G divided by c to some power. Let's just say that it's missing factors of c . I may comment on that in some notes that I will add to the board. And so that kind of tells you that this is going to be a really small quantity. Makes it very hard for these things to be measured. Nonetheless, it turns out even though they have-- if you think about the impact they have on the detector, even though it's small, they carry a tremendous amount of energy.

I'm going to take a break before I begin recording my next lecture. But the next lecture is actually going to lead us to an understanding of a different version of the quadrupole formula. It's based on this, but it describes the energy content of gravitational waves. Both of these formulas go by the name the quadrupole formula. It's a little bit confusing sometimes if both are used. And you're going to use both of them on upcoming homework assignments. All right. So I will end this recording here.