PROFESSOR: OK. Let's resume. I cut things off at a time when we had the final answer. And I left you hanging because we don't know what the final answer is telling us. This says that as we change the orientation of the normal to the plate, relative to x3 and, secondarily, as we change the angle between the threefold axis and x2, we get this strange third-rank trigonometric function.

Let's convert these cosines into the appropriate angles. This was the twofold axis. And this was x1. This was the direction of the threefold axis. This was x3. And this is the direction x2. And that comes out in between a pair of twofold axes.

And we may want to look at this from up above, relative to these twofold axes that occur. This thing here looks like a trigonometric identity, doesn't it? Let's let this angle here be theta. And l1, then, is cosine of theta.

And so our geometry is that this is l1 cosine of theta. This is l2, which is the cosine of the angle between our direction and x2. And that is the cosine of pi over 2, minus theta. And that, then, is equal to sine of theta.

So this identity, as you all know-- I'm not telling you anything that you don't already know-- this is d is equal to d1,1 1, times the cosine of 3 theta, right? You knew that. Believe it or not, it is. It's one of these obscure trigonometric identities, which wouldn't occur to you in a million years unless you go digging through some handbooks.

So this is rather astonishingly simple. It's simply a cosine function. But the interesting thing is it goes as a function of 3 theta. So this goes through one cycle between a pair of adjacent twofold axes.

So it starts out as d1,1. And it finishes up at d1,1. And in between it is minus d1,1.

And so it goes down.

And now we come to something that I do differently than most people do. Most people will say it goes as minus d1,1. But minus d1,1 I like to show as a lobe going off in this direction, with a minus sign.

AUDIENCE: How do you get that cosine of 3 theta?

PROFESSOR: Hmm?

AUDIENCE: How do you get d equals d1 cosine 3 theta?

PROFESSOR: How do I get that?

AUDIENCE: Yeah.

PROFESSOR: That is just a trigonometric identify, believe it or not, a well-known trigonometric identify. Actually, it's an exceedingly obscure trigonometric identity.

So the way this is going to go from x1 is it's going to be a positive lobe around x1. There's going to be a negative lobe along x2 and then a positive lobe again about the twofold axis that's 120 degrees away. And then a negative lobe, and then a positive lobe opposite this negative lobe, and a negative lobe opposite this positive lobe. So it's a six-membered-- six lobes.

There is always a positive lobe opposite a negative lobe. And what this means is that the charge is of opposite sign on opposite ends of the twofold axes. So the response peaks up on twofold axes.

Now the thing that I don't like is that what Nye does is to say, OK, this is a negative value. So you should plot the radius in a negative direction. And that puts it over here, right on top of this positive lobe.

So what Nye shows is the polar plot of this result, in the plane of the twofold axes, is simply this. This is x1. And he shows a lobe here. And he shows a lobe here. And he shows a lobe here.

And if you interpret that as a polar plot you say, well, I know what the value is in here. It's decreasing. And there's nothing going on in here. So the response must be 0.

And then it starts coming up again. And there's a response in different directions here. And then it goes back down to 0. And in this range, there's nothing going on.

And that's not true. What's going on is a negative value of the modulus. And to me, that becomes abundantly clear if you just put a sign that labels the sign of the modulus within those lobes. Now I submit, that's pretty anisotropic, isn't it?

Yes, Steve?

AUDIENCE: Did you just arbitrarily choose where your positive and negatives go, is it [INAUDIBLE] out here?

PROFESSOR: No. This comes out here. When l1 is 0, that's-- excuse me. When the angle is 0, l1 is plus 1. Remember that l1 is the cosine of this angle. When that angle is 0, then the value is plus 1. So I was careful to put the label x1 on this.

Now the other thing that we should examine is how these lobes vary in a direction perpendicular to x, in a direction that includes x3. So let us look at how the function varies in a direction perpendicular to-- that includes x1 and x3. OK. Call this angle phi. And in the x1, x3 plane, d1,1,1 prime turns out to be l1 cubed d1,1. Before you go to our general expression.

And l2 is cosine of 90. This thing drops out. We're left with simply d1,1 prime equals l cubed times d1,1. So this then goes as cosine cubed of phi, which is a lobe that starts out at plus 1. And then because it's cosine cubed, this dies out very, very rapidly. So these lobes are very flat in the x1, x2 plane, die out very rapidly as the cube of phi, where phi is the angle between the twofold axis and x3.

So the interesting thing about this surface is that if you decided to pick up the random fragment of crystal and determine what its Piezoelectric Modulus is. One way of doing this, a poor man's test for piezoelectricity is to just clamp a fragment of crystal between two electrodes and then hook this up to a variable frequency generator that sweeps through a range of frequencies, changes the frequency of a voltage across these plates, and then comes back and sweeps again. Or alternatively, have a knob one on that lets you change frequency. Then put a pair of earphones across the crystal.

And if we would do that, by having this variable frequency generator, as we turned it, when we hit a resonant frequency that set up either a full wavelength of a vibrational wave in the crystal-- or half wavelength, or one wavelength, or 3/2 wavelength-- there'd be a resonance. And the capacitance does something crazy. It does like this as you go through the resonant frequency. And then your simple detector-- a pair of earphones-- as you tuned the frequency, you would hear static as you went through this discontinuity in capacitance.

Or alternatively, you could put this on a cathode ray tube and just have a sweep frequency and put the voltage across the electrodes on the oscilloscope screen. And you would see something like this, and then this discontinuity, and then maybe second harmonic. And it would work just fine. This is a poor man's way of detecting piezoelectricity. But what if-- what if-- you happen to have your piece of quartz with the electrodes directly on the c-axis? The modulus would be 0, and you wouldn't detect anything.

Or if you put on the probes such that they were in a direction that was in between these lobes where, again, the Piezoelectric Modulus has gone to 0. You wouldn't find anything. So measuring the Piezoelectric Modulus for a random chunk of material is dangerous, if you just look at one direction and say nothing's going on; it's not piezoelectric. The other thing is the material could have a piezoelectric response that's so weak you just can't measure it. Yeah, OK?

AUDIENCE: This stuff's only for single crystals, right? Because if you [INAUDIBLE] polycrystal material, you're getting [INAUDIBLE] averaging.

PROFESSOR: You're averaging over all directions. So this is for a single crystal. And that's the only time you get these exotic, very anisotropic surfaces.

Another method that I've seen in a rather old book is really nice. What you do is you cut your crystal into a little plate. And then you drive the plate by means of putting electrodes on it and hook this up to a variable frequency.

And if the material is transparent, you will set up-- depending on the velocity of sound in the material-- you will set up standing waves when you hit the right frequency. And these lines that I've drawn are places where the displacement is always 0. And these things that I've indicated at little squares is a region where the displacement goes up and down between its maximum and minimum extreme.

Now if this crystal is transparent and you shine a light through it, these little regions bounded by lines of zero displacement act like little lenses. So if you pass a beam of light through it, you get a bunch of right maxima, little focused spots on the sheet. And you can determine for a particular frequency if this is the 1 and 1/2 wavelengths along a dimension that you know, you can again find the Piezoelectric Moduli. So there are lots of ways of detecting this effect. But when you have a single crystal, you've got to really look at this as a function of direction, to be absolutely sure.

One other Piezoelectric Modulus is in the problem set. 4 bar 3m, this is the structure of sphalerite and a lot of the compound semiconductors. This is not in the notes. But if you try the problem, what you find is, again, a highly, highly anisotropic behavior. And this is the direction of the unit cell of the crystal. You find that the piezoelectric response is a series of very sharp lobes, positive and negative, going along the directions of the body diagonal, namely the 1, 1, 1 directions in the crystal.

And then there's another lobe that goes like this, negative, positive. Another lobe that goes up here, positive and negative. So very sharp lobes along the body diagonal, alternately positive, negative, positive, negative as you go around 0, 0, 1 plane. OK. That's what the longitudinal piezoelectric effect is like.

You can invent single-crystal devices for weighing fish. And then you have an interesting question. If you hang the weight on the crystal, and it's a flat plate, how would you orient the crystal to get the optimum sensitivity? So you define a scalar modulus for the crystal in the particular direction you're exerting a uniaxial stress. And then you express this module in terms of theta, the angle of rotation, within the crystographic plane.

In the case of the sup cell if you looked at that geometry, that is a flat plate subject to compression, where you measure the surface charge perpendicular to the direction of the uniaxial stress, not parallel to it as we've done here. And then there's a question if you rotate the plane of the crystal about the imposed stress, what orientation gives you optimum, maximum response? So there are fun things that you could do with that.

We have a quiz next time. There's a lame-duck period after the quiz will be over. And I want to say a little bit about the tensor aspects of elastic moduli. Mechanical behavior is something that you have or will cover in great detail in a graduate-level class on mechanics, but I think very few people who deal with things other than cubic materials or with single-crystal materials where the elastic constants are functions of crystal symmetry and direction. So let me, at least at the time available, set up a definition of stress in terms of strain.

Stress is a second-rank tensor. There are six unique components. So we could write an expression for sigma 1,1; sigma 2,2; sigma 3,3; sigma 4,4; and sigma 5,5. And sigma--

AUDIENCE: --1,2.

PROFESSOR: No, I want to go down like this. So this would be sigma 2,3; sigma 1,3; and sigma 1,2; eventually to be known as sigma 1, sigma 2, sigma 3, sigma 4, sigma 5, and sigma six. But actually, if we're writing out a full tensor relation, we should write down all six elements here. So let me do that and put in a 3,2. Put in all nine of them, because before we condense to matrix form, that is what we're going to have.

So we'll have a sigma 2,3; we'll have a sigma 3,2; we'll have a sigma 1,3; a sigma 3,1; a sigma 1,2; and a sigma 2,1. And we can express each of the stresses in terms of the elements of strain. So I'll have an epsilon 1,1; an epsilon 1,2; an epsilon 2,2; an epsilon 3,3; an epsilon 2,3; an epsilon 3,2; an epsilon 1,3; an epsilon 3,1; an epsilon 1,2; and an epsilon 2,1. Nine terms, nine terms.

And in between is going to be a tensor consisting of 9 by 9 elements. There will be 81 different coefficients in here. If we want to derive symmetry constraints, we're going to have to, for every one of those 81 elements, do a transformation. And each transformed element is going to be a sum of four direction cosines times one tensor element, repeated 81 different times. Not something to be undertaken on a short afternoon.

OK. What are the tensor elements in here? They are represented by the symbol c. This would be c1,1,1,1. This would be c1,1,2,2 times epsilon 2,2; a c1,1,3,3; a c1,1,2,3; and so on. And these c's, in one of the great perversions of scientific notation, are called stiffnesses.

The other thing we could do is to write strain in terms of stress. And really, stress is something you can do to the crystal. So stress is an independent variable.

And this is something that I feel more at home with. So I'll have an epsilon 1,1; 2,2; 3,3; epsilon 3,2; 2,3; epsilon 1,3; epsilon 3,1; epsilon 2,1; and an epsilon 1,2. Again, nine elements of strain. And these will be given by coefficients times each of the nine elements of stress, sigma 1,1; sigma 2,2; sigma 3,3; sigma 2,3; sigma 3,2; and so on for nine different components.

Andy the coefficients here are designated with the symbol s. These are the tensor elements. So this is s1,1,1. This is s1,1,2,2. And the s's stand for compliances.

Now English is a very strange language. But to call compliance s and to call stiffness c is surely a perversion that can be designed for no other purpose than to confuse the introductory student and give the instructor some feeling of superiority. I don't feel superior, I feel embarrassed that I have to explain this. Stiffness, c; compliance, s. It's got be one of the atrocities of scientific notation.

To remember which goes with which, I'll tell you what works for me. If it doesn't work for you, forget it. The s goes with sigma. And the little c goes with the epsilon. That's a c with a bar in it. That works for me. If it doesn't work for you, use your own mnemonic device. English does a lot of things strangely, but nothing is as perverse as this.

Now what's some examples? Consider this. You park your car in your driveway, but you drive your car on a parkway. Why don't you drive your car on a driveway and park your car in a parkway? No, it's the other way around. It's almost as bad as this.

What are some other ones? In my basement, I have something that's called a hot water heater. You don't heat hot water, you heat cold water. Why isn't it called a cold water heater?

Even better, I once worked for a year and a half in Switzerland. And in the kitchen of my home I had, in German, an elektrowarmwassersheissungsapparat. That's one word that's 11 syllables. The Germans have a knack with the language that we have never approached in English.

But there's another one you've probably heard. You know what the prefix "pro" means. That means for, and "con" means against. And so progress is moving forward, and I leave it to you to decide what Congress means.

OK. So I could go on and on. I probably shouldn't, unless you egg me on. No, I won't. But anyway, there are-- OK, I'll give another one.

Why is what a doctor does called "practice"? This is not reassuring at all. Let me invite you over to my practice.

And why is it when you take an airplane your flight ends with a terminal? That's doesn't sound very encouraging either. So anyway, language is silly.

And I guess if you want to call c stiffnesses and s compliances, it's OK. The verbal description of these terms means something. Because if the stiffness is very high, that says that a little bit of strain requires that you haven't posed a very large stress. So stiff is when the material is resistant to stress. So you have to have a very, very large strain to get yourself a given level of stress.

Conversely, if the material is very compliant, a very small stress for a compliant material should give you a big strain. And that's exactly what a large value of s will do. It will give you a large strain for a relatively small stress. So the words describe what's going on.

OK. There is great utility in reducing these equations to a matrix form, which is going to cut us down from a 9 by 9, or a tensor with 81 elements, to a smaller number of subscripts and a smaller number of matrix elements, so if we let sigma 1 be sigma 1,1; if we let sigma 2 be equal to sigma 2,2 and sigma 3 be equal to sigma 3. And then if we write the inequalities we've had before, and we let epsilon 1,1 to be replaced by epsilon 1, then we would have out in front here a term simply c1,1, a 1 for the sigma and a 1 for the epsilon.

The next term would be c1,2 times epsilon 2, where obviously epsilon 2 has replaced epsilon 2,2, and 1 stands for the two indices, 1,1. So this will go c1,3; epsilon 3. And then you hit these messy factors of two. You'll have a pair of terms. You'll have c1,4. And that will stand for a term c1,1,2,3. And then there'll be another term c1,1,3,2.

And if you want to write this as a matrix, c1,4 has to be equal to the sum of these two things, times epsilon 4. That's the only way you can write it as a matrix. And if you think it was bad worrying about a factor of 2 in front of half of the terms on the right-hand side of the equation, we're going to be dealing with factors of 2 for terms in the upper 3 by 3 array. We're going to be worrying about a factor of 4, somehow or other, down in the lower 3 by 3 array. So the way you handle the factor of 2 really is a nightmare for relations in elasticity.

But continuing on with the first line, c1,5 times epsilon 5 would be a combination of c1,1,3,1 plus 1,1,1,3. And then finally, when you get to c1,6 times epsilon 6, that would represent a combination of c1,1,2,1 plus c1,1,1,2. You have to make a decision where to swallow the 2. And probably the best thing I can do is to give you conventions for relabeling stress, strain, stiffness, and compliance, which I'll pass around to you.

No big deal, except that this is a case where you can eat the 2's early on, but you can't even break even. So very, very briefly, what you do is you convert the tensor elements of stress to matrix elements in exactly the same way we did with piezoelectricity. For strains, we do exactly the same thing that we did with the converse piezoelectric effect.

The factor of 2 that you swallowed in definition of stress pops up to haunt you when you deal with the strains. And you cannot write a nice matrix relation unless you divide the off-diagonal elements of strain by 1/2. And that was exactly the same thing we encountered with the converse piezoelectric effect, which related strain to an applied electric field.

In defining the matrix stiffnesses, the c's, you forge straight ahead, and you let ci,j,k,l, which is identical to ci,j,l,k-- that's the term that relates the two equal shear stresses-- and you define that as a matrix term with two subscripts. No factors of 2 or 4 are involved. Then when you hit the s's, there's a nightmare. si,j,k,l is sm,n if m and n are 1, 2, or 3; that is to say, not 4, 5, or 6.

si,j,k,l, which is the same as s,i,j,l,k, has to be defined as one half of sm,n, where m or n is 4, 5, or 6. And then finally, you have to throw in a factor of 4 when both m and n are 4, 5, or six or, in other words, m and n are not 1, 2, or 3. So it is a mess. And these are rules that you have to bear in mind if you're ever going to go from matrix form, which works absolutely lovely in a fixed coordinate system. But if you have a single crystal and you want to refer it to different axes, you have to be prepared to resurrect the full tensor form. And then and only then can you work symmetry transformations.

I don't want to go through the simple algebra of expanding and contracting these terms. So the next two sheets show you how you condense from tensor to matrix notation, and then go back from matrix notation to tensor notation. This is for the compliances si,j,k.

And the next page does the same thing, if you care to write stress in terms of strain.

And there are a lot more factors of 2 that appear there. But again, it shows you how you go from tensor notation, for subscripts on the c's, down to matrix notation with 2 subscripts, and then go back up again to tensor notation. So it's a tedious business. And you've got to keep careful tracks of your factors of 2 and your factors of 4.

How about symmetry restrictions? Holy mackerel. Transformation of 81 different elements. Well, in order to do the tensor transformations, you have to go to the full-force subscript notation. Only then is a law of transformation defined. Fourth-rank properties are even tensors.

So mercifully, there are not as many different possibilities. All point groups that differ by presence or absence of inversion have exactly the same form of the tensor. So symmetry 2, m, and 2/m look alike. Symmetry 2, 2mm, and 2/m, 2/m, 2/m look alike. So there's only one orthorhombic tensor, only one kind of monoclinic tensor, only one kind of triclinic tensor, and so on.

And the only place that you have more than one form of the tensor for a particular set of point groups is looking at the point groups that are based on single-rotation axes, something like 4, 4 bar 4m, and those that are based on the axial arrangements 4, 2, 2. So they're different, just as they were for second-rank tensor properties.

For cubic, a rather remarkable result. There are three independent compliances. But cubic crystals are not elastically isotropic. They are anisotropic. And it's just the nature of the tensor that requires that. So rather than letting you hang by your thumbs wondering what's on these pages, let me pass around the summary of symmetry constraints.

If you understood what we did for third-rank tensors or even second-rank tensors, you how to do this. And fortunately, the terms almost over. So I can't make you do it, which is probably an enormous relief to you.

Some additional bits of information. The transformations for hexagonal crystals are complicated because threefold and sixfold axes do not change the reference axes x1, x2, x3 into one another. And the equations that relate the individual tensor elements are more complex.

For the threefold axis, for example, 2 s1,1 minus 2 s1,2 is equal to s4,4. So they're not simple relations between them. And that's simply because you're not changing one reference access into another. You're changing x1, for example, into a linear combination of x1 and x2.

Two things down here of interest. I said that cubic crystals are not isotropic. What would have to be the case if the material were to be isotropic? Well, it turns out that if you do this for the compliances, s, if 1/2 of s4,4 is equal to s1,1 minus s1,2, then the material is elastically isotropic. If you do this in terms stiffnesses, the c's, then 2 c4,4 is equal to c1,1 minus c1,2 if the material is supposed to be isotropic or going to be isotropic.

Then there is one other inequality that depends on structure. And this is something known as the Cauchy relation. And it depends on the interatomic forces. The first condition that has to be met is that the forces between the atoms should be central forces.

What is a central force? Well, this is a case where the attractive force between the atoms is directly along the line joining their centers. Isn't that always the case? Don't crystals hold together because there's an attractive force between atoms? Well actually, for metallic crystals and ionic crystals, maybe that's not a bad assumption, particularly for ionic crystals.

But if you had a covalent material, where the bonding was due to overlap of orbitals like this, then the thing that holds the crystal together is overlap between these orbitals. And the force holding the atoms together is a force that goes through this shared electron pair. And that's not a central force. And so the Cauchy relation generally fails rather badly for covalently bonded materials.

Second assumption is that each atom is at a center of symmetry. And the reason for that is so the force in the plus-x direction is the same as the force in the minus-x direction. And that is not true for any material that has tetrahedral coordination. So it's not true for any of the forms of SiO2. It's not true for any of the compound semiconductors that are based on tetrahedral units. So there, the atom inside of these tetrahedra is decidedly not at an inversion center.

And finally, you assume that the crystal is under no state of initial stress. Because if it is under initial stress, you've squished it, you've stretched the bonds. And the forces-- not to a major degree, to be sure, but-- the forces will not truly be perfect central forces. So if all three of these assumptions are satisfied, then you have an additional equality c1,2 is equal to c4,4.

What I'll do next time is I'll bring in some examples of data for stiffnesses and compliances and, in particular, show you how some of these elastic tensor elements change with temperature. And we can examine them to see how well the Cauchy equality or the isotropy condition is satisfied.

All right. That's about all that we'll do about the definition of fourth-rank tensor properties. On next Tuesday to wrap the term up, we'll look at the variation of some elastic moduli, such as Young's Modulus or the shear modulus, for different symmetries as a function of direction in the crystal, OK?