TITLE: Global Symmetry, the Phonon Density of States and RUS

TITLE: Proceedings of the Resonance Meeting. Volume 1. Transcripts

GLOBAL SYMMETRY, THE PHONON DENSITY OF STATES AND RUS

T.W. DARLING AND A. MIGLIORI
LOS ALAMOS NATIONAL LABORATORY

ABSTRACT

RUS measures (in principle) a complete set of the acoustic phonons near the Brillouin zone center. By comparison with a model of an homogeneous solid of known global symmetry and judicious application of a goodness-of-fit measure, the elastic constants can be deduced.

For many solid state phase transitions the change in the symmetry of the object locally and globally is of prime importance, particularly where the transition leads to disorder at some length scale. We will discuss the application of RUS to determination of symmetry and whether in practice this can provide useful information on systems with fine or meso-scale structure.

Such systems are in some sense "intermediate" between homogeneous single crystals (low symmetry) and an homogeneous, isotropic ("glassy") state (high symmetry). We expect this to be reflected in the behavior of the phonon density-of-states and thus into other physical properties such as the specific heat. We discuss the application of RUS to measuring this behavior directly or to guiding other experimental techniques which have access to different length scales.

These topics will be illustrated by analyses of our RUS data on perovskite and martensite systems.

TRANSCRIPT

DR. DARLING: When I started to write this talk a week ago, based loosely on an abstract submitted 2 months ago, I thought I was going to have a lot to say about the thermodynamics that you can extract from RUS measurements and, in fact, what I came up with was a bunch of reminders of undergraduate physics and a set of object lessons, cautionary object lessons, which are actually kind of relevant to some of the talks we have seen.

[Transparency 1]

The first one was relative size. In Douglas Adams' work, *The Hitchhiker's Guide to the Galaxy*, one of the worst fates imaginable is to be put in the total perspective vortex where your ego is demolished by a view of your own size in relationship to the universe.

[Transparency 2]

Here is the universe of the vibrations of a solid and down here on a microscopic dot is where we work and if we think that we are going to extrapolate something we measure down here into the amount of detail that exists over the whole phonon spectrum, then we have got to be pretty careful about what we do here.

At this end of the universe, this scale, which is accessible by scattering experiments, neutrons, x-rays, various optical techniques, strange variations of Mössbauer measurements, generally tend to cut off at about 10^{11} Hz. Our measurements down here tend to cut off usually below 8 MHz of resonances in a system for a number of reasons.

Up here we have -- generally harmonic approximation is used to calculate these modes and that gets rid of a whole load of experimental useless pieces of things like thermal expansion and temperature dependencies of elastic constants, so the calculations up here are fairly straightforward, but the models used actually rely on the fact that the dissipation processes at high frequencies are limited to, say, electronic components and harmonic parts, whereas down here, of course, we have a whole bunch of other things.

High frequencies: We measure total energies by replacing the solid by a whole bunch of harmonic oscillators.

[Transparency 3]

This gives us amplitudes and frequencies for oscillators. You then put in wave numbers by applying a model that gives you plane waves inside the solid, but you end up throwing away any dissipation processes and all of these useless parts and, quite often, you deal with it in a Debye model where you are modeling a continuum.

[Transparenciey4]

Down at our end of the world we have dissipation processes, and I am not going to talk very much about this, just to say this looks like one of the curves that Bob Leisure measures for relaxation processes, which are all low-frequency items.

We have trouble in general dealing with this because of this frequency variation across a scan and only if you measure the temperature dependence with respect to a single peak do you actually get information that you can really extract sensibly.

[Transparency 5]

And we have this problem. This is a tantalum sample and this is a set of plots of a 20-kHz region that I have taken from about 1 MHz up to 4 MHz and, as we can see down here, the density of states is low enough that we can easily select our peaks, but perhaps by the time we have reached 2.5 MHz we are getting to the point where we can no longer extract peaks out of this thing usefully and as we go farther and farther up we find we get into this limit where the density of states is high enough that the phases and amplitudes are averaging out and we are seeing no signal, not because the electronics is incapable of it but just because the material itself is cutting us off, so the 8 MHz that I showed on the first plot is perhaps an upper limit.

[Transparency 6]

How well do we actually do with densities of states? Here is the first cautionary object lesson. This is a strontium titanate sample, a single crystal, it was cut and we measured the green triangles to show the density of states up to about 8 MHz, perhaps a little farther.

It follows exactly -- by accident here, the scaling factor is actually one -- a curve, dN/df follows a curve of f^2. The object lesson involved here are the red and the black dots, which are the results of fits from the program that most of us used, which was supplied by Bill Visscher and various variants that people work on here. It fails, quite seriously, here, and this is not surprising when you look at it, because we know that in order to make the basis functions for this fit we divide up into 8 different parity-organized states and we have another parameter that describes each mode, which is essentially harmonic numbers, which are k and i in the fitting program.

If we are using basis functions that are only going up to an order of 10 in a polynomial, that is about the number of nodes that we expect to see, so in an harmonic idea of waves inside a solid, we expect to see only up to 10 nodes inside this, so when we get to values of i, that is, harmonics that are greater than 10, we are not going to be fitting this material, we do not have a basis function that is capable of actually following it.

What this means -- and this is the i value here, which is approximately 10 and perhaps some variants of the program (we have 12 order polynomials) -- is the total number of lines you are going to measure to get there is perhaps somewhere between 50 and 60 lines. This will vary a little bit depending upon the geometry of the sample and the actual relative sizes of the C_{ij}'s, but it means that if you are going to measure more than 50 lines and fit it, you are not really doing yourself any favors at all; all you are doing is introducing errors.

DR. ISAAK: What shape is the specimen?

DR. DARLING: It is a rectangular parallelepiped with an aspect ratio of perhaps 2:1 on its large face and 4:1 -- it is a plate. You will not be able to measure more than the order of the polynomial times 8 lines, because then you will almost certainly have exceeded the capability of the basis functions to describe the vibration of the solid.

[Transparency 7]

Of the thermodynamic things we can measure, specific heat is something that is usually looked at. This, of course, ignores all the electronic components. You can describe it by the density of states, multiply it by the specific heat per oscillator, integrate it over the whole of accessible space to give you -- this is in the harmonic approximation as well.

This function here, which, if you use an Einstein specific heat per oscillator, you can see cuts off the integral as the temperature gets lower (there is actually an invisible yellow line there that goes almost straight across at 20 K).

If we are living in this little space of the density of states here, then the only time that we really have anything useful to say about specific heat accurately is when we are down at temperatures of a few micro-Kelvin, which few of us routinely approach. Otherwise, every single thing that we measure is a Debye solid, and that is all we can extract from the density of states out of the RUS system, so we lose, there is no way we will see any of the detail at higher frequencies.

[Transparency 8]

The same thing largely goes for the dispersion curves that can be measured, once again, by neutron-scattering experiments. These are just examples from the literature for silicon and nickel. The blue line is basically the lower limit of experimental data from scattering experiments. The little red bit here is round about where we live in these curves.

What I have really put these up to show is that this is considered to be pretty good agreement between data and experimental fits for calculating via various lattice dynamic approaches to the dispersion curves. In fact, they are not really very good and, in particular, the only way that these types of approaches have anything to say about the zone center phonons, which is where we are, is by the theoretical extrapolations coming down to the gamma point at the zone center.

This is kind of interesting. Several years ago a publication came out looking at the properties of gold zinc and, based upon data that fitted about the same as this looks, they calculated elastic constants that were a good 200% wrong and published a statement saying that the gold zinc, because of the elastic constants, the low anisotropy from their measured elastic constants, would never show a martensitic transformation, which is now known to be untrue.

[Transparency 9]

This is the measurement of the frequencies of a single crystal of gold zinc and, in fact, it is well known to have a martensitic transformation.

The errors that are introduced by extrapolating high-frequency measurements back down to where we live can be significant. It is very important, in fact, that we make good measurements of the dispersion curves of the zone center to keep these guys honest, to make sure that their models are not actually erring off into strange places.

This also shows something that Don brought up the other day, which is an enormous temperature dependence -- the transformation actually takes place at about 60 K, but up to room temperature we still have negative elastic constant variations with temperature. One of the things I want to do is to actually extend this to see just how far it goes; it shows no tendency to be turning over even that far away from the transformation, and I will mention a little more about that later.

[Transparency 10]

Every result that we derive from our RUS measurements is only as good as the Debye model and this can matter a lot at intermediate temperatures between sort of helium and nitrogen temperatures, where the actual shape of density-of- states curves and the turnover of dispersion curves can matter.

The Debye model is very popular. It is essentially a single parameter model where we extract a sort of geometric average velocity from our measured transverse and longitudinal wave velocities, which we have extracted from our measurements of c_{ij}.

The thing to note, really, about this is that to do this sensibly we really have to know how many independent branches there are of the phonon spectrum, and it means that we have to be very sure of the number of elastic constants that we are using to fit our spectrum and that, of course, is the global symmetry of the sample.

[Transparency 11]

Once again, using a sample of strontium titanate -- this is just to show that we actually measured strontium titanate -- I took a set of these room-temperature spectra and, using the program, again, to fit, I produced, with the same input data, I fitted a variety of different symmetries. I know it is cubic but I would like the program to be able to tell me it is cubic.

One of the first things you notice is the RMS number that comes out of the program actually can hide a lot of problems. Because it is a percentage error it normalizes the size of the error with the frequency, so as the frequency goes up, whatever the error is, it starts shrinking as 1 over the frequency, so at very high frequencies its contribution to the error measure is small.

I had a look at the actual chi^2 per degree of freedom, which you expect to be around 1 for a good fit, and I was able, just on the basis of that, to reject isotropic and hexagonal fits. In the second line, the iso-star line, I took only 4 lines, the first 4 lines of this measurement, and ran a fit and I got a wonderful RMS value out of the program, but it is a rubbish fit. It has no relationship to the real elastic constants, so a small number of lines can be a serious problem here.

The 3 hexagonal fits are just the 3 possible orientations of the c axis within these samples, supposing I did not know it, but I can reject all of these 32 lines (I can reject the pieces with isotropic components). The cubic one fits pretty well, it has a good chi^2 degree of freedom, good RMS error, but when you get up to tetragonal and orthogonal, I expect these to fit well. The cubic is a special case of tetragonal and orthogonal and, in fact, I expect the fit to improve slightly when I do use these symmetries because I have more free parameters to fit with.

The interesting thing is that in tetragonal, for some reason in this system, I have 2 -- these 2 (the green and the blue are absolutely invisible on this) -- the 2 and 3 fit from the tetragonal are actually the ones that I expect. They are a little bit better than the cubic fit and they have the same values of the elastic constants and the bulk modulus.

But there is one orientation of this that has an excellent fit but it is wrong, it has provided the wrong elastic constants, so if I were to do this, if I were just to start it off with a tetragonal assumption, I might say, great, I have one orientation that fits really well and gives me some values, and two orientations that are in the wrong c-axis direction, so I might simply throw away cubic as the symmetry of this material just because I seem to have a consistent tetragonal fit.

I have an orthogonal fit as well, which looks like the wrong tetragonal fit and I need to be able to reject these possibilities when I am making these measurements; otherwise I will end up with the wrong number of elastic constants, the wrong symmetry and all sorts of other wrong things when I try to do physics.

The thing that you really need to look at here is the chi^2 error, the actual function, the chi^2 function, at the minimum. In fact, you need to be able to reject fits on the basis of those things as well, so if you are just reading the RMS error out of the program, you are doing yourself a disservice, because you need to be also rejecting places that have too shallow a fit in any of the elastic constants at the minimum.

That means the numbers that are produced by the program (and these are the numbers that come out of the program), the 2% chi^2 confidence limit, it is not really an error, it is a confidence limit that we have placed in the chi^2 hyper-surface. If these numbers exceed a few percent, 2 or 3%, you need to reject them as well, because the fit is too shallow for you to be certain that there is actually not another minimum in there somewhere farther down the road that has not been reached by the program.

This is the second cautionary tale, that you have to be very careful upon reading just this value of the RMS error, because you can reach false minima that look excellent in terms of that error but unless they also have sharp minima they do need to be rejected.

PARTICIPANT: Could you have sharp minima that were also false?

DR. DARLING: I do not know. I have asked myself that question a lot, as to whether or not you can achieve a sharp minimum that is not the right one. I have never seen it, actually, but that would be a wonderful question to be able to answer. I do not know in general if you can do it in an n dimensional space with a chi$_2$ hyper-surface and say whether the global minimum has always got to be sharp and that no other minimum will also be sharp.

[Transparency 12]

There is a set of magic numbers that we need to take note of and I call them "magic" because they are largely the result of hand-waving and experience but maybe ignoring them is sometimes useful, too.

We know that the number of lines had better be less than the order of the polynomial times the number of parity states, because, otherwise, you will be trying to fit rubbish lines. I have known people who have said, "Hey, I've measured 120 lines and got this incredible fit," and you know that the top 50 lines just do not have any meaning and that the RMS error, because it is a fractional error, has just hidden the fact that there are huge problems there.

In fact, it may turn out to be much lower if you end up using up -- if you have a very low elastic constant you might use up a whole load of harmonics associated with that type of parity motion before you get to 80 lines. By the time you have reached $i = 8$ for any of the parity modes, you have started to exceed the program's capability to describe the vibration.

You need to use a sufficient number of lines. We arbitrarily suggest that the number of elastic constants times 5 is a good number of lines to be using, at least, because fewer lines will also give you false fits. The reason the 4 lines gave me a wonderful fit with strontium titanate is that the first 3 depended on only one elastic constant, as Don showed us earlier, so 4 lines I can easily fit with 2 parameters to an isotropic fit, because 3 of the lines depend on only one elastic constant, and that is a very common thing with RUS samples.

When you are looking at the errors coming out of the program, we arbitrarily say that if we do not get better than 0.5%, then we do not really have a fit. We would really like it to be much better than that, in fact, better than 0.2%, but the criterion is that you also, if you are going to determine the symmetry of the sample, really have to know how many elastic constants belong to it, you need to make sure that the confidence limit that we have applied does not exceed, let's say, 2 or perhaps 3%.

When you, in fact, increase the number of free parameters and you get an improvement in the RMS error between, say, cubic and tetragonal, numbers of less than, say, 20%, it may not actually have any meaning for it at all just because of the improvement of the number of free parameters.

[Transparency 13]

We are making a Debye model. If we want to calculate free energies, entropies, any of these things that we are extracting numbers from based on Debye temperatures, then we have to understand that we are operating in a Debye model that has, up until the zone edge, an f^2 dependence of the density of states and a linear dependence of the dispersion of the phonons all the way to the zone edge.

If anything depends at intermediate temperatures on the fine structure, that is, cutoffs of the various longitudinal and transverse modes, then we are going to run into problems. We are only as good as the Debye model.

However, measurements, in general, are so good that we definitely need to do them in order to keep the models honest that are measured from high-frequency things. The symmetry is really very important and we determine only the global symmetry and there are various scaling things (a bunch of slides I did not show you that say that we actually miss a whole bunch of interesting regions), where micro-structure and grain structures start to make the wave vector a poor quantum number.

We can see sometimes things to do with micro-structures by looking at the way the symmetry is changed, so, for instance, a polycrystal will look isotropic as opposed to being the symmetry of each single crystal, and you can perhaps extract some things. One of the things I did not talk about very much, the temperature dependence of C_{ij}, which has a connection to the average oscillator amplitude of each of these normal modes, may eventually play back onto a lot of these things. There may be some things you can extract from temperature dependences and we do do temperature dependences very well in RUS, particularly the anomalous varieties of it, such as Don showed and such as gold zinc shows, to be able to correlate that with such things as anomalous thermal expansions and to extract that from our measurements would be, I think, a useful thing and, I think, a profitable way to look.

Thanks.

DR. ANDERSON: I always approach this problem as saying if I am going to measure a single crystal, I have got a good single crystal, so I am going to do x-rays and find out what this crystal is. Then I will know what the symmetry is and then I will put that into the inverse calculation. I never dreamed that you would do it the other way.

DR. DARLING: There is a problem there in that quite a lot of x-ray determinations will not show you small distortions or small deviations if you look at your average Laue back-reflection pattern. The difference between a slightly tetragonal distorted material and a cubic one is very small and you would likely say this is cubic.

We are somewhat more sensitive in measuring the vibrations. We could actually see if there was a small tetragonal distortion and that is kind of important and, also, if we go through a phase transformation, you do not generally get the chance to be watching with x-rays, not many people have the equipment to do that, but we get to see if the symmetry has changed going through a phase transformation, and we could determine the symmetry, in principle, after having gone through a phase transformation.

DR. ANDERSON: What you are saying is that we have the power through RUS determination to confer the symmetry that you thought you had from an x-ray and you might be surprised.

DR. DARLING: Absolutely. You might be surprised.

DR. MIGLIORI: Take an x-ray of a pattern of a quasi-crystal and that tells you what to do.

DR. DARLING: So you will see a fivefold pattern out of a quasi-crystal but, elastically, it is extraordinarily isotropic.

DR. ISAAK: I think you have highlighted it. It is still pretty dicey to say you can determine the symmetry. You showed those slides and just from my own experience, I have done that, too, where I have taken like rutile, which is tetragonal, I tried an orthorhombic fit (I was just checking some things) and I came into similar problems like that. I was really surprised.

When you give it extra degrees of freedom, you take those liberties and go off somewhere -

DR. DARLING: You need to be extraordinarily careful to do that and sometimes you do not know exactly what the symmetry is. There are some people at the lab who, for reasons of their own, like to splat little rods of aluminum into a brick wall so that it plastically squashes, and I had a look at some of these pieces and I am pretty sure that there is at one unique axis, which is the one perpendicular to the direction that splashed this piece of aluminum, but it turns out that the cross-section you get is elliptical as well, which suggests that the perpendicular axes can be special.

It was very interesting to them to know exactly what the symmetry elastically was to our best measurements of that material, whether, in fact, it was orthorhombic, as the shape suggested -- we, in fact, found that the difference between tetragonal and orthorhombic, elastically, there was very small.

DR. ANDERSON: If you put in a certain symmetry the machine will find, your program will find a minimum for you and you will say "aha."

DR. DARLING: Absolutely, you need to be watching, and that is why the question of whether a false minimum can actually also be sharp, can actually fit all of these criteria that I put down, is a very interesting one, but I do not know the answer to it and I have not seen a false minimum that is sharp, but it would be very interesting to know whether you can be led truly astray. We know that the solution is not unique in principle.

DR. DEMAREST: When you are calculating the low-temperature specific heat, you can probably get even better results if you do not use the RUS calculation but, rather, these periodic boundary conditions.

DR. DARLING: That is certainly true. One of the problems with RUS is that the free boundary conditions mix up the modes. In periodic boundary conditions you manage to keep everything essentially as plane waves that fit inside the particular solution. With free boundary conditions, as everybody knows, your modes are never pure anything. The picture of the squashed cubes that is so popularly shown just tells you there are no plane waves living in the RUS spectrum until you get to high-enough frequencies. That is certainly a complication for calculations of specific heats.

If you needed to, you could, in fact, use that spectrum to get detailed real specific heats out of the real modes at the low-frequency end of the spectrum.

DR. LEDBETTER: Your results imply that you can take one crystal and measure the RUS spectra and get information about whether there are zone-center soft modes?

DR. DARLING: A zone-center soft mode is usually an optical -- I have not looked at zone-center acoustical soft modes which are not a result of effects at other parts of the zone. Certainly, if you have a zone-edge soft mode as in strontium titanate, we see the result here because of the pulling down of the whole dispersion curve, so we see an effect in the gradient and the dispersion at the zone center, but if such a thing exists as an acoustical soft mode only at the zone center we would see it.

DR. ANDERSON: But metals do not have optic modes and they have transitions. so --

DR. DARLING: Gold zinc is a metal and it has got optical modes. You just need to have more than one atom per primitive cell to do that.

DR. MIGLIORI: Orson, a first-phase order transition can go without a soft mode.

DR. ANDERSON: Okay, yes.

DR. DARLING: Thank you.