

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high quality educational resources for free. To make a donation, or view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

PROFESSOR: Alright. So what we are doing is covering 8.334, which is statistical physics. AND let me remind you that basically statistical physics is a bridge from microscopic to macroscopic perspectives. And I'm going to emphasize a lot on changing perspectives.

So at the level of the micro, you have the microstate that is characterized, maybe, by a collection of momenta and coordinates of particles, such as particles of gas in this room. Maybe those particles have spins. If you are absorbing things on a surface, you may have variables that denote the occupation-- whether a particle or site is occupied or not. And when you're looking at the microscopic world, typically you are interested in how things change as a function of time. You have a kind of dynamics that is governed by some kind of a [INAUDIBLE] which is dependent on the microstate that you are looking at, and tells you about the evolution of the microstate.

At the other extreme, you look at the world around you and the macro world. You're dealing with things that are described by completely different things. For example, if you're thinking about the gas, you have the pressure. You have the volume. So somehow, the coordinates managed to define and macrostate, which is characterized by a few parameters in equilibrium.

If you have something like a collection of spins, maybe you will have a magnetization. And again, another important thing that characterizes the equilibrium we describe as temperature. And when you're looking at things from the perspective of macro systems in equilibrium, then you have the laws of thermodynamics that govern constraints that are placed on these variables. So totally different perspectives. And what we have is that we need to build a bridge from one to the

other and this bridge is provided by statistical mechanics.

And the way that we described it in the previous course is that to what you need is a probabilistic prescription. So rather than following the time evolution of all of these degrees of freedom, you'll have a probability assigned to the different microstates that are dependent on the macro state. And for example, if you are dealing in the canonical ensemble, then you know you're at the particular temperature this form has E to the minus beta, this Hamiltonian that we you have here governing the dynamics. Beta was one over KT . So I could write it in this fashion.

And the thing that enabled us to make a connection between this deterministic perspective and this equilibrium description with probabilities abilities was relying on the limit where the number of degrees of freedom was very large. And this very large number of degrees of freedom enabled us to, although we had something that was probably realistic, to really make very precise statements about what was happening.

Now, when we were doing this program in 8-333, we looked at very simple systems that were essentially non interacting, like the ideal gas. Or we put a little bit of weak perturbation and by some manipulations, we got things like liquids. But the important thing is that where is this program we could carry out precisely when we had no interactions.

In the presence of interactions we encountered, even in our simplified perspective, new things. We went from the gas state to a liquid state. We didn't discuss it, but we certainly said a few things about solids. And clearly there are much more interesting things that can happen when you have interactions. You could have other phases of matter that we didn't discuss, such as liquid crystals, superconductors, and many, many other things. So the key idea is that you can solve things that are not interacting going their own ways, but when you put interactions, you get interesting collective behaviors.

And what we want to do in 8-334, as opposed to just building the machinery in 8-333, is to think about all of the different types of collective behavior that are possible

and how to describe them in the realm of classical systems. So we won't go much into the realm of quantum systems, which is also quite interesting as far as its own collective behaviors is concerned. So that's the program.

Now, what I would like to do is to go to the description of the organization of the class, and hopefully this web page will come back online. So repeating what I was saying before, the idea is that interactions give rise to a variety of interesting collective behaviors, starting from pretty much simple degrees of freedom such as atoms and molecules and adding a little bit of interactions and complexity to them, you could get a variety of, for example, liquid crystal phases.

But one of things to think about is that you think about all of the different atoms and molecules that you can put together, and the different interactions that you could put among them-- you could even imagine things that you construct in the lab that didn't exist in nature-- and you put them together, but the types of new behavior that you encounter is not that much. At some level, you learn about the three phases of matter, gas, liquids, solids, of course, as I said, there are many more, but there are not hundreds of them. There are just a few.

So the question is why, given the complexity of interactions that you can have at the microscopic scale, you put things together, and you really don't get that many variety of things at the macroscopic level. And the answer has to do with mathematical consistency, and has at its heart something that you already saw in 8-333, which is the central limit theorem. You take many random variables of whatever ilk and you add them together, and the sum has a nice simple Gaussian distribution. So somehow mathematics forces things to become simplified when you put many of them together.

So the same thing is happening when you put lots of interacting pieces together. New collective behaviors emerge, but because of the simplification, in the same sense the type of simplification that we see with the central limit theorem, there aren't that many consistent mathematical descriptions that can emerge. Of course, there are nice new ones, and the question is how to describe them. So this issue of

what are possible consistent mathematical forms is what we will address through constructing these statistical field theories.

So one of the important things that I will try to impose upon you is to change your perspective. In the same way that there is a big change of perspective thinking about the microscopic and macroscopic world, there is also a change of perspective involved in the idea of starting from interacting degrees of freedom, and changing perspective and constructing a statistical future. And I'll give you one example of that today that you're hopefully all familiar with, but it shows very much how this change of perspective works, and that's the kind of metallurgy that we will apply in this course.

Basically the syllabus is as follows. So initially I will try to emphasize to you how, by looking at things, by averaging over many degrees of freedom, at long length scales and time scales, you get simplified statistical field theory descriptions. The simplest one of these that appears in many different contexts is this Landau-Ginzburg model that will occupy us for quite a few lectures. Now once you construct a description of one of these statistical field theories, the question is how do you solve it.

And there are a number of approaches that we will follow such as mean field theory, et cetera. And what we'll find is that those descriptions fail in certain dimensions. And then to do better than that, you have to rely on things such as perturbation theory. So this is kind of perturbation theory that builds upon the types of perturbation theory that we did in the previous semester, what is closer to the kind of perturbation theory that you would be doing in quantum field theories.

Alongside these continuum theories, I will also develop some lattice models, that in certain limits, are simpler and admit either numerical approaches or exact solutions. The key idea that we will try to learn about is that of Kadanoff's perspective of renormalization group and how you can, in some sense, change your perspective continuously by looking at how a system would look over larger and larger length scales, and how the mathematical description that you have for the system changes as a function of the scale of observation, and hopefully becomes simple at

sufficiently large scales.

So we will conclude by looking at a variety of applications of the methodologies that we have developed. For example, in the context of two dimensional films, and potentially if we have time, a number of other systems. The first thing that I'm going - the rest of the lecture today has essentially nothing in common with the material that we will start to cover as of the second lecture. But it illustrates this change in perspective that is essential to the way of thinking about the material in this course. And the context that I will use to introduce this perspective is phonons and elasticity.

So you look around you, there's a whole bunch of different solids. There's metals, there's wood, et cetera. And for each one of them, you can ask things about their thermodynamic properties, heat content, for example, heat capacity. They are constructed of various very different materials. And so let's try to think, if you're going to start from this left side from the microscopic perspective, how we would approach the problem of the heat content of these solids. So you would say that solids, if I were to go to 0 temperature before I put any heat into them, they would be perfect crystals.

What does that mean? It means that if I look at the positions of these atoms or particles that are making up the crystal-- I guess I have to be more specific since a metal is composed of nuclei and electrons-- let's imagine that we look at the positions of the ions. Then in the perfect position they will form a lattice where I will pick three integers, L M N and three unit vectors, and I can list the positions of all of these ions in the crystal. So this combination L M N that indicates the location of some particular ion in the perfect solid that's indicated by, let's say, the vector r . So this is the position that I would have ideally for the ion at zero temperature that is labelled by r .

Now, of course when we go to finite temperatures, the particles, rather than forming this nice lattice, let's imagine a square lattice in two dimension, starts to move around. And they're no longer going to be perfect positions. And these distortions we can indicate by some vector u . So when we have perfect crystals across

deformations, this Q star changes to a new position, Q of r , which is its ideal position plus a distortion field u , at each location. OK?

Now associated with these change of positions, you have moved away from the lowest energy configuration. You have put energy in the system. And you would say that the energy of the system is going to be composed of the following parts.

There's always going to be some kind of a kinetic energy. That's sum over r p^2 squared over $2m$. Let's imagine particles all have the same mass m . And then the reason that these particles form this perfect crystal was presumably because of the overlap of electronic wave functions, et cetera. Eventually, if you're looking at these coordinates, there's some kind of a many body potential as a function of the positions of all of these particles.

Now, what we are doing is we are looking at distortions that are small. So let's imagine that you haven't gone to such high temperatures where this crystal completely melts and disappears and gives us something else. But we have small changes from what we had at zero temperature. So presumably the crystal corresponds to the configuration that gives us the minimum energy. And then if I make a distortion and expand this potential in the distortion field, the lowest order term proportional to various u 's will disappear because I'm expanding around the minimum.

And the first thing that I will get is from the second order term. So I have one half sum over the different positions. I have to do the second derivative of the potential with respect to these deformations. Of course, if I am in three dimensions, I also have three spatial indices, x , y , and z , so I would have to take derivatives with respect to the different coordinates, α and β , and summed over them. And then I have u_α of r , u_β of r prime. And then, of course, I will have higher order terms in this expansion. This is a general potential, so then in higher order terms, it would be order of u cubed and higher. OK? Fine?

So I have a system of this form. Now, typically, the next stage if I stop at the quadratic level-- this I would do for a molecule also, not only for a solid-- is to try to

find the normal modes of the system. Normal modes I have to obtain by diagonalizing this matrix of second derivatives. Now, there are a few things that I know that help me, and one of the things that I know is that because the original structure was a perfect solid, let's say, then there will be a matrix-- sorry, there will be an element of the second derivative that corresponds to that r and this r prime.

That's going to be the same as the second derivative that connects these two points, because this pair of points is obtained by the first pair of points by simple translation among the lattice. The environment for this is exactly the same as the environment for this. So essentially, what I'm stating is that this function does not depend on r and r prime separately, but only on the difference between r and r prime. So, I know a lot-- so this is not if I had n atoms in my system, this is not something like n squared over 2 independent things, it is a much lower number. The fact that I have such a lower number allows me to calculate the normal modes of the system by Fourier transform.

I won't be very precise about how we perform Fourier transforms. Basically I start with this k alpha beta, which is a function of separation. I can do a sum over all of these separations r of e to the $i K \cdot r$, for appropriately chosen k . And summing over all pairs of differences, so the argument r here now, is what was previously r minus r prime. So basically, what I can do is I can pick one point and go and look at all of the separations from that point. Construct this object, this will give me the Fourier transformed object that depends on the vague number k .

So if I look at the potential energy of this system minus its value at zero temperature, which from one perspective was one half sum over r r prime alpha beta, is k alpha beta, r minus r prime, u alpha of r , u beta of r prime, in the quadratic approximation. If I do Fourier transforms what happens because it is only a function of r minus r prime, and not r and r prime separately, is that in Fourier space, it separates out into a sum that depends only on individual K modes.

There's no coupling between k and k prime. So here we have r and r prime, but by the time we get here, we just have one k , alpha and beta. We'll do one example of

that in more detail later on. I have the Fourier transformed object, and then I have $u_{\alpha k}$ Fourier transform. So in the same manner that I Fourier transformed this kernel, $k_{\alpha\beta}$, I can put here a u and end up here with a \tilde{U} and \tilde{u}_{β} of k^* .

So we start over here, if you like. If I have n particles with a matrix that is n by n -- actually $3n$ by $3n$ if I account for the three different orientations-- and by going to Fourier transforms, we have separated it out for each of n potential Fourier components, we just have a three by three matrix. And so then we can potentially diagonalize this three by three matrix to get two eigenvalues, λ_{α} of K . Once I have the eigenvalues of the system, then I can find frequencies or eigenfrequencies. ω_{α} of K , which would be related to this λ_{α} of K divided by m .

So you go through this entire process. The idea was you start with different solids. You want to know what the heat content of the solid is. You have to make various approximations to even think about the normal modes.

You can see that you have to figure out what this kernel of the interaction is-- Fourier transform it, diagonalize it, et cetera And ultimately thing that you're after is that there are these frequencies as a function of wave number. Actually, it's really a wave vector, because there will be three different-- K_x , K_y and K_z . And at each one of these K values, you will have three eigenfrequencies. And presumably as you span K , you will have a whole bunch of lines that will correspond to the variations of these frequencies as a function of K .

Why is that useful? Well, the reason that it's useful is that as you go to high temperature, you put the energies into these normal modes and frequencies. That's why this whole lattice is vibrating.

And the amount of energy that you have put at temperature T on top of this V_0 that you had at zero temperature, up to constants of proportionality that I don't want to bother with is a sum over all of these normal modes that are characterized by K and α , the polarization and the wave vector. And the amount of energy, then, that

you would put in one harmonic oscillator of frequency ω . And that is something that's we know to be $\hbar \omega \alpha$ of K , divided by e to the beta $\hbar \omega \alpha$ of K minus 1.

So the temperature dependence then appears in this factor of beta over here. So the energy content went there, and if you want to, for example, ultimately calculate heat capacity, I have to calculate this whole quantity as a function of temperature. And then take the derivative.

So it seems like, OK, I have to do this for every single solid, whether it's copper, aluminium, wood, or whatever. I have to figure out what these frequencies are, what's the energy content in each frequency. And it seems like a complicated engineering problem, if you like.

Is that it about this that transcends having to look at all of these details and come to this? And of course, you know the answer already, which is that if I go to sufficiently low temperature, I know that the heat capacity due to phonons for all solids goes like t^3 . So somehow, all of this complexity, if I go to low enough temperature, disappears. And some universal law emerges that is completely independent of all of these details-- microscopics, interactions, et cetera.

So our task-- this is the change of perspective-- is to find a way to circumvent all of these things and get immediately to the heart of the matter, the part that is independent of the details. Not that the details are irrelevant. Because after all, if you want to give some material that functions at some particular set of temperatures, you would need to know much more than this t^3 law that I'm telling you about.

But maybe from the perspective of what I was saying before-- how many independent forms there are, in the same sense that adding up random variables always gives you a Gaussian. Of course, you don't know where the mean and the variance of the Gaussian is, but you are sure that it's a Gaussian form. So similarly, there is some universality in the knowledge that, no matter how complicated the material is, its low temperature heat capacity is t^3 . Can we get that by an

approach that circumvents the details?

So I'm going to do that. But before, since I did a little bit of hand-waving, to be more precise, let's do the one-dimensional example in a little bit more detail. So my one-dimensional solid is going to be a bunch of ions or molecules or whatever, whose zero temperature positions is uniformly separated by some lattice spacing, a , around one dimension.

And if I look at the formations, I'm going to indicate them by the one-dimensional distortion U_n of the n th one along this chain. So then I would say, OK, the potential energy of this system minus whatever it is at zero, just because of the distortion, I will write as follows-- it is a sum over n . And one thing that I can do is to say that if I look at two of these things that are originally at distance a , and then they go and the deform by U_n and U_{n+1} , the additional deformation from a is actually U_{n+1} minus U_n . So I can put some kind of Hookean elasticity and write it in this fashion.

Now of course, there could be an interaction that goes to second neighbours. So I can write that as $\frac{K_2}{2} (U_{n+2} - U_n)^2$ and third neighbors, and so forth. I can add as many of these as I like to make it as general as possible.

So in some sense, this is a kind of rewriting of the form that I had written over here, where these things that were a function of the separation-- these $K_{\alpha\beta}$ of separation or these K_1, K_2, K_3 , et cetera-- in this series that you would write down. Now if you go to Fourier space, what you can do is each U_n of U_n , is the distortion in the original perspective, you can Fourier transform. And write it as a sum over k to the ik position the n th particle is na , times \tilde{u}_k .

And once you make this Fourier transform in the expression over here, you get an expression for $V - V_0$ in terms of the Fourier modes. So rather than having an expression in terms of the amplitudes u_n , after Fourier transform, I will have an expression in terms of \tilde{u}_k . So let's see what that is. Forget about various proportionality.

I have the sum over n . Each one of the U_n 's I can write in this fashion in terms of U

tilde of K . Since this is a quadratic form, I need to have two of these. So I will have the sum of k and k prime.

I have the factor of $1/2$. Each U_n goes with a factor of e to the $i n a k$. But then I had $2 U_n$. There's a term here, if I do the expansion, which is U_n squared. So I have one from k and one from k prime.

However, if I have $U_n + 1 - U_n$, what I have is e to the $i k a$ minus 1 . I already took the contribution that was e to the $i n k$. From the second factor, I will get e to the $i k$ prime a minus 1 .

This multiplies K_1 over 2 . And then I will have something that's K_2 over 2 , e to the $2 i k a$ minus 1 , e to $2 i k$ prime a minus 1 , and so forth. Multiplying at the end of the day U tilde of k , U tilde of k prime.

Now when I do the sum over n , and the only n dependence appears over here, then this is the thing that forces k and k prime to add up to 0 , because if they don't add up to 0 , then I'm adding lots of random phases together. And the answer will be 0 . So essentially, this sum will give me a delta function that forces k plus k prime to be 0 .

And so then the additional potential energy that you have because of the distortion ends up being proportional to $1/2$, sum over the different k 's. Only one k will remain, because k prime is forced to be minus k . And so I have U of k , U of minus k , which is the same thing as U of k complex conjugate. So I will get that.

And then from here, I will get K_1 . Now then, k prime is set to minus k . And when I multiply these two factors, I will get $1 + 1 - e$ to the $i k a$ minus e to the minus $i k a$. So I will get $2 - 2 \cosine$ of $k a$. And then I will have K_2 , $2 - 2 \cosine$ of $2 k a$, and so forth.

Why are the lights not on?

OK, still visible. So yes?

AUDIENCE: So in your last slide, you have an absolute value U_k , but wouldn't it be-- right above it, is it U of k times U star of k prime, or how does that work?

PROFESSOR: OK, so the way that I have written is each U_n I have written in terms of U tilde of k . And at this first stage, the two factors of U_n that I have here, I treat them completely equivalently with indices k and k prime. So there is no complex conjugation involved here. But when k prime is set to be minus k , then I additionally realize that if I Fourier transform here, I will find that U tilde of minus k is the same thing as U tilde of k star. Because essentially, the complex conjugation appears over here.

It's not that important a point. The important point is that we now have an expression for our frequencies, or ω alpha of k -- actually, there's no polarization. It's just ω of k -- are related to square root of something like a mass down here. Again, that's not particularly important, but something like k_1^2 , 2 minus 2 cosine of ka , k_2^2 , 2 minus 2 cosine of $2ka$, and so forth.

So I can plot these frequencies ω as a function of k . One thing to note is first of all, the expression is clearly symmetric under k goes to minus k , so it only depends on cosine of k . So it is sufficient to draw one side. The other side for negative k would be the opposite.

The other thing to note is that again, if I do this Fourier transformation, and I have things that are spaced by a , it effectively means that the shortest wavelength that I have to deal with are of the order of a , which means that the wave numbers are also kind of limited by something that I can't go beyond. So there's something that in the generalized over here, you recall that your k vectors are within the Brillouin zone. In one dimension, the Brillouin zone is simply between minus π/a and π/a .

Now the interesting thing to note is that as k goes to 0, ω goes to 0. Because all of these factors you can see, as k goes to 0, vanish. In particular, if I start expanding around k close to zero, what I find is that all of these things are quadratic. They go like k^2 . So when I take the square root, they will have an absolute value of k .

So I know for sure that these omegas start with that. What I don't know, since I have no idea what k_1 , k_2 , et cetera are, is what they do out here. So there could be some kind of a strange spaghetti or going on over here, I have no idea. There's all kinds of complexity. But they are away from the k close to 0 part.

Again, why does it go 0? Of course, k equals to 0 corresponds to taking the entire chain and translating it. And clearly, I constructed this such that all of the U 's are the same-- I take everything and translate it-- there's no energy cost. So there's no energy costs for k equals to 0.

The energy costs for small k have to be small. By symmetry, they have to be quadratic in case, so I take the square root and we will get a linear. Of course, you know that this linear part-- we can say that omega is something like a sum velocity. So all of these chains, when I go to low enough case or low enough frequencies, admit these sound-like waves.

Now heat content-- what am I supposed to do? I'm supposed to take these frequencies, put them in the expression that I have over here, and calculate what's going on. So again, if I want to look at the entirety of everything that is going on here, I would need to know the details of k_2 , k_3 , k_4 , et cetera. And I don't know all of that. So you would say I haven't really found anything universal yet.

But if I look at one of these functions, and plot it as a function of the frequency, what do I see? Well, omega goes to 0. I can expand this.

What I get is kt . Essentially, it's a statement that low frequencies behave like classical oscillators. A classic oscillator has an energy kt . Once I get to a frequency that is of the order of kt over \hbar , then because of the exponential, I kind of drop down to 0.

So very approximately, I can imagine that this is a function that is kind of like a step. It is either 1 or 0. And the change from 1 to 0 occurs at the frequency that is related to temperatures by k_2 over a .

So if I'm at some high temperature, up here, and I want to-- so this omega is $k v t$ sum i over $h \bar{}$ -- that's the corresponding high frequency-- I need to know all of these frequencies to know what's going on for the energy content. But if I go to lower and lower temperatures, eventually I will get to low enough temperatures where the only thing that I will see is this linear portion. And I'm guaranteed that I will see that. I know that I will see that eventually.

And therefore, I know that eventually, if I go to low enough temperatures, the excitation energy becomes low enough, it's simply proportional to this integral from 0. I can change the upper part of the infinity if I like. $dk h \bar{ } 1k$ divided by e to the $\beta h \bar{ } dk$ minus 1,

And again, dimensionally, I have two factors of k here. Each k scales with kt . So I know that whole thing is proportional to kt squared.

In fact, there's some proportionality constants that depend on $h \bar{ } , t, et cetera$. It doesn't matter. The point is this t squared.

So I know immediately that my heat capacity is proportional to-- derivative of this is going to be proportional to t . The heat capacity of a linear chain independent of what you do. So no matter what the state of interactions is, if I start with a situation such as this at 0 temperature, I know if I put energy into it at low enough temperature, I would get this heat capacity that is linear.

I don't know how low I have to do. Because how I have to go to depends on what this velocity is, what the other complication is, et cetera. So that's the part that I don't know. I know for sure of the functional form, but I don't know the amplitude of that functional form.

So the question is, can we somehow get this answer in a slightly different way, without going through all of these things? And the idea is to do a coarse grain. So what's going on here? Why is it that I got is form?

Well, the reason I got this form was I got to low enough temperature. At low enough temperature, I have only the possibility of exciting emote, whose frequencies were

small. I find that frequencies small correspond to wave numbers k that are small, or they correspond to wavelengths that are very large.

So essentially, if you have you solid, you go to low enough temperature, you will be exciting modes of some characteristic wavelength that are inversely proportional to temperature, and become larger and larger. So eventually, these long wavelength modes will encompass whole bunches of your atoms. So this λ becomes much larger than the spacing of the particles in the chain that you were looking at.

And what you're looking at. low temperature, is a collective behavior that encompasses lots of particles moving collectively and together. And again, because of some kind of averaging that is going on over here, you don't really care about the interactions among small particles. So it's the same idea. It's the same large n limit appearing in a different context. It's not the n that becomes very large, but n that becomes of the order of, let's say, 100 lattice spacings-- already much larger than an individual atom doing something, because it's a collection of atoms that are moving together.

So what I drew here was an example of a mode. But I can imagine that I have some kind of a distortion in my system. Now, I started with the distortions U_n that were defined at the level of each individual atom, or molecule, or variable that I have over here. But I know that things that are next to each other are more or less moving together. So what I can do is I can average. I can sort of pick a distance-- let's call it the x -- and average of all of those U_n 's that are within that distance and find how that averages.

And as I move my interval that I'm averaging, I'm constructing this core function U of x . So there is a moving window along the chain constructed with the V_x which is much larger than a , but is much less than this characteristic frequency. And using that, I can construct a distortion field.

I started discrete variables. And I ended up with a continuous function. So this is an example of a statistical field. So this distortion appears to be defined continuously.

But in fact, it has much less degrees of freedom, if you like, compared to all of the discrete variables that I started with. Because this continuous function certainly does not have, when I fully transform it, variations at short length scales.

So we are going to be constructing a lot of these coarse grained statistical fields. If you think about the temperature in this room, bearing from one location to another location pressure, so that we don't strike sound waves, et cetera. All of these things are examples of a continuous field.

But clearly, that continuous field comes from averaging things that exist at the microscopic level. So it's kind of counter-intuitive that I start with discrete variables, and I can replace them with some continuous function. But again, the emphasis is that this continuous function has a limited set of available and surveyed numbers over which it is defined. OK.

So we are going to describe the system in terms of this. So that analog of this potential that we have over here is some b function of this U of x . And I want to construct that function.

And so the next step after you decided what your statistical field is to construct some relevant thing, such as an potential energy that is appropriate with that statistical field, putting as limited amount of information as possible in construction of that. So what are the things that we are going to put in constructing this function at.

The first thing that I will do is I will assume that there is a kind of locality. By which, I mean the following. While this is in principle the function of the entire function, locality means that I will write it as an integral of some density, where the density at location x that I'm integrating depends on you at that location. But not just you, also including derivatives of you.

And you can see that this is really a continuum version of what I have written here, this. If I go to the continuum, this goes like a derivative. And if I look at further and further distances, I can construct higher and higher derivatives.

So in the sense that this is a quite general description, I can construct any kind of potential in here by choosing interactions. K_1 , K_2 , K_3 , K_{100} that go further and further apart, you would say that if I include sufficiently high derivatives here, I can also include interactions that are extending over far away distances. The idea of locality is that while you make this expansion, our hope is that at the end of the day, we can terminate this expansion without needing to go to many, many higher orders.

So locality is two parts. One, that you will write it in this form. And secondly, that this function will not depend on many, many high derivatives.

The second part of it is symmetries. Now, one of the things that I constructed in here and ultimately was very relevant to the result that I had was that if I take a distortion U of x and I add a constant to everybody, so if I replace all of my U s to plus U plus 5, for example, the energy does not change. So V of this is the same thing as V of U of x .

So that's a symmetry. Essentially, it's this translation of symmetry that I was saying right here at the beginning, that this only depends on the separation of two points. It's the same thing. But what that means is that when you write your density function, then the density cannot depend on U of x .

Because that would violate this. So you can only be start with things that depend on the U of V IBX this second, et cetera. So this is 1. This is 2.

Another thing is what I call stability. You are looking at distortions around a state that corresponds to being a stable configuration of your system. What that means is that you cannot have any pairs in this expansion that are linear.

So again, this was implicit in everything that we did over here. We went to second order. But there's third order, et cetera, are not ruled out. It is more than that.

Because you require the second order terms to have the right sign, so that your system corresponds to being at the bottom of a quadratic potential rather than the top of it. So there is a bit more than the absence of linear terms. So given that, you

would say that your potential for this system as a function now of this distortion is something like an integral over x .

And the first thing that is consistent with everything we have written so far is that if we be proportional to du by dx squared. So there's a coefficient that I can put here. Let's call it k over 2 . It cannot depend on you. It has to be quadratic function of derivative.

That's the first thing I can write down. I can certainly write down something like d^2u by dx to the fourth power. And if I consider higher order terms, why not something like the second derivative squared, first derivative squared, a whole bunch of other things. So again, there's still many, many, many terms that I can't write down. Yes?

Is that second term supposed to be a second derivative to the fourth power?

PROFESSOR: Yes. Thank you. So that when I fully transform this, the quadratic part becomes sum over k , k over 2 . This fully transformed, becomes k squared. This fully transformed, as you said, is second derivative squared.

So it becomes k to the fourth. I have a whole bunch of terms. And then I have U of k squared. And then I will have higher order terms from Fourier fully transform of this. Yes?

AUDIENCE: Does this actually forbid odd derivatives? Are you saying the third derivative and stuff don't--

PROFESSOR: I didn't go into that, because that depends on some additional considerations, whether or not you have a mirror symmetry. If you have a mirror symmetry, you cannot have terms of that are odd and x . Whether or not you have some condition on U and minus U may or may not forbid third order terms in $[U]$ $[U]$ by the x .

So once I go beyond the quadratic level, I need to rely on some additional symmetry statement as to which additional terms I am allowed to write down. Yes?

AUDIENCE: Also the coefficients could depend on x , right?

PROFESSOR: OK. So one of the things that I assumed was this symmetry, which means that every position in the crystal is the same as any other position. So here, if I break and make the coefficient to be different from here, different from there, it amounts to the same thing, that the starting point was not the crystal.

AUDIENCE: Shouldn't that be written as-- in the place where you wrote down the symmetry, it should be U of x plus c inside the parenthesis?

PROFESSOR: No. No. So look at this. So if I take U_n and I replace U_n to U_n plus 5, essentially I take the entire lattice and move it by a distance. Actually, 5 was probably not good. 5.14. It's not just anything I can put over here. Energy will not change. OK?

AUDIENCE: That must be different from adding to all the ends a constant displacement.

PROFESSOR: Ends are labels of your variables. So I don't know what mean by--

AUDIENCE: OK, you're right. But in that picture, where instead of N 's, we have X ?

PROFESSOR: Yes.

AUDIENCE: It seems like displacing in space would mean adding up to x .

PROFESSOR: No. No. It is this displacement. I take U_1, U_2, U_3, U_4 . U_1 becomes U_1 plus .3. U_2 becomes U_2 plus .3. Everybody moves in step.

AUDIENCE: So the conclusion is the coefficients don't depend on x ?

PROFESSOR: If you have a system that is uniform-- so this statement here actually depends on uniformity. This is an additional thing, uniform. So one part of the material is the same. Now, you have non-uniform systems.

So you take your crystal and you bombard it with neutrons or whatever. Then you have defects all over the place. Then one location will be different from another location. You are not able to write that anymore. So uniformity is another symmetry that I kind of implicitly used. Yes?

AUDIENCE: When your uniform is a separate part, why isn't it implied by translational symmetry?

PROFESSOR: If I take this material that I neutron bombarded, and I translate it in space, it's internal energy will still not change, right?

AUDIENCE: OK. OK.

PROFESSOR: So again, once I come to this stage, what it amounts to is that I have constructed a kind of energy as a function of a deformation field, which in the limit of very long wavelengths has this very simple form which is the integral du by dx squared. There are higher order terms. But hopefully in the limit of long wavelengths, the higher derivatives will disappear.

In the limit of small deformation, the higher order terms will disappear. So the lowest order term at long wavelengths, et cetera, is parametrized by this $1/k$. If I fully transport it, I will just get k over $2k$ squared. When I take the frequency that corresponds to that, I will get that kind of behavior.

So by kind of relying on these kinds of statements about symmetry, et cetera, I was able to guess that. Now, let's go and do with this for the case of material in three dimensions. Actually, in any dimensions, in higher dimensions.

So I take a solid in three dimensions, or maybe a film in 2 dimensions. I can still describe it's deformations at low enough temperature in terms of long wavelength modes. I do the coarse graining. I have U of x . Actually, both U and x are now vectors.

And what I want to do is to construct a potential that corresponds to this end. OK? So I will use the idea of locality. And I write it as an integral over however many dimensions I have. So d is the dimensionality of space of some kind of an energy density.

And the energy density now will depend on U . Actually, U has many components, U_α of x . Derivatives of U -- so I will have du_α by dx_β . And you can see that higher order derivatives, they're really more and more indices. So the complication now is that I have additional indices involved.

The symmetry that I will use is a slightly more complicated version of what I had before. I will take the crystal, U of x , and I can translate it just as I did before. And I just say that this translation of crystal does not change its energy.

But you know something else? I can take my crystal, and I can rotate it. The internal energy should not change. So there is a rotation of this that also you can put the c inside of. Before or after rotation, it doesn't matter. The energy should not depend on that. OK.

So let's see what we have to construct. I can write down the answer. So first of all, we know immediately that the energy cannot depend on U itself for the same reason as before. It can depend on derivatives. But this rotation and derivatives is a little bit strange.

So I'm going to use a trick. If I do this in Fourier space, like I did over here, I went from this to k over 2 integral dk k squared U tilde of k squared. If I stick with sufficiently low derivatives, only at the level of the second order derivatives, if I have a second order form that depends on something like this, I can still go to Fourier space. And the answer will be of the form integral d dk .

The different k modes will only get coupled from higher order terms, third order terms, et cetera. At the level of quadratic, I know that the answer is proportional to d dk . And for all of the reasons that we have been discussing so far, the answer is going to be U of k squared times some function of k , like k squared, k to the fourth.

Now, whatever I put over here has to be invariant under rotations. So, let's see. I know that the answer that I write here should be quadratic [INAUDIBLE] tilde. It should be at least quadratic in k , because I'm looking at derivatives. In the same way that I had k here, I should have factors of k here.

But k is a vector when I go to three dimensions. U becomes a vector when I go to three dimensions. So I want to construct something that involves quadratic in vector k , quadratic in vector U , and is also invariant under rotations.

One thing that I know is that if I do a dot product of two vectors, that dot product is invariant on the rotations. So I have two vectors. So I know, therefore, that $k \cdot k$ is a rotational invariant. The square of k is rotationally invariant.

But also, $k \cdot U$ square of k is rotationally invariant. OK? So what I can do is I can say that the most general form that I will write down will allow 2 terms. The coefficients of that are traditionally called μ . This is $\mu/2$. This one is called $\mu + \lambda/2$. Actually, I have to put an absolute value squared here.

So that's the most general theory of elasticity in any number of dimensions that is consistent with this symmetry that I have here. And it turns out that this corresponds to elasticity of materials that are isotropic. And they are described by 2 elastic coefficients, μ and λ , that are called [INAUDIBLE] coefficient.

μ is also related to shear modulus. Actually, μ and λ combined are related to [INAUDIBLE] And if I want, in fact, to fully transform this back to the space, in real space the can be written as $\mu/2 \sum_{\alpha\beta} u_{\alpha\beta}^2$, where the sum over α and β takes place. α and β run from 1 to d . And the other term, $\lambda/2$, is $U_{\alpha\alpha}$ of x , $U_{\beta\beta}$ of x . And this object, $U_{\alpha\beta}$, is one half of the symmetrized derivatives, $du_{\alpha} / dx_{\beta} + du_{\beta} / dx_{\alpha}$. And it's called the strength

AUDIENCE: Question.

PROFESSOR: Yes?

AUDIENCE: So are you still looking in the regime of low energy expectations?

PROFESSOR: Yes. That's right.

AUDIENCE: So wouldn't the discreteness of the allowed wave vectors become important? And if so, why are you integrating rather than discrete sum?

PROFESSOR: OK. So let's go back to what we have over here. The discreteness is present over here. And what I am looking at for the discreteness, this spacing that I have

between these objects is 2π over L . The L is the size of the system.

So if you like, what we are looking at here is the hierarchy of landscapes, where L is much larger than the typical wavelengths of these excitations that are set by the temperature, which is turn much larger than the lattice place. And so when we are talking about, say, a solid at around 100 degrees temperature or so, this-- then say, over here, it typically spans 10 to 100 atoms, where as the actual size of the system spans billions of atoms or more. And so the separations that are imposed by the discreteness of k are irrelevant to the considerations that we have. Yes?

AUDIENCE: So before, with this adding a constant c , that corresponds to translating almost crystal by some vector.

PROFESSOR: Right.

AUDIENCE: For the rotation, is this a rotation of the crystal or is this a rotation of the displacement field?

PROFESSOR: It's the rotation of the entire crystal. So you can see that essentially both x and U have to be rotated together. I didn't write it precisely enough. But when I wrote the invariant as being $k \cdot U$, the implicit thing was that the debate vector and the distortion are rotated.

AUDIENCE: So does it require an isotropic crystal in that case?

PROFESSOR: Yes.

AUDIENCE: I would think if you're rotating everything together, who cares if one axis is different than another? Because if I have a non-isotropic crystal and I rotate it around, it shouldn't change the internal energy.

PROFESSOR: OK. Where it will make difference is at higher order terms. And so then I have to think about the invariants that are possible at the level of higher order terms. But that's a good question. Let me come back and try answer that more carefully next time around.