PROFESSOR: It's sort of a little bit thrown in between. What you've done is you've essentially just finished doing the whole section on energy models. Remember we started all the way in the beginning with empirical potential models when doing [INAUDIBLE] method, and then Professor Marzari did quantum mechanics.

A lot of the rest of what you're going to see is on particular simulation techniques that bring in dynamics and temperature into the problem. What I wanted to do with this lecture is sort of tell you the story of my life.

You work so hard on doing quantum mechanics, and you do it well. And you come up with all these beautiful things you calculate in energies, and you give a whole talk on that.

And there'll always be some smarty pants in the audience that will get up at the end of your talk and say, you know. But all your calculations are at 0 temperature. And so basically with one slight comment, they basically wipe your whole talk off the table.

And half of the room suddenly sort of thinks that you're just kind of a funny curiosity. So what I want to do here today is teach you sort of to what extent does temperature matter. So first of all, when are zero-temperature calculations relevant?

When are they extremely good to mimic the real world, which is non-zero temperature? When is it so so? And when do you absolutely need to take temperature explicitly into account? And then the next step is, how do you do that?

Now I will give you hints of that today already-- how you do that. But later the rest of the course is essentially on that. So to learn to push the right buttons here. Huh? So again these are the objectives of today.

And maybe what I'll add a little is sort of, why is it so hard to include temperature? People who don't do simulations often think that, why don't you just turn on that button? There's this button temperature. And why don't you turn it on?

Why don't you just do zero-temperature calculation? So but if you look at the Kohn-Sham equations, obviously, there is no temperature in it. OK? So this is not like some physical approximation that we can decide to turn on or off. It turns out if you want to deal with temperature effects, there's a whole lot of work you have to do on top of your energy calculation.

You can see why that is. Essentially what you're doing with quantum mechanics is you're calculating the ground state energy. Now essentially temperature is no more-- when you're at non-zero temperature, it's exactly making the same statement in saying, you know, I've put energy in my system to make it be above the ground state energy.

So at non-zero temperature, you've pumped in extra energy. There's many ways you can see that. You can just see that from the derivative of the energy with temperature, which is the heat capacity, which is always positive, which means as soon as you raise the temperature from 0, the energy has gone up above your ground state energy.

So where is that energy? This is going to be the key question when you want to deal with temperature simulation. First of all, let me tell you there is no generic temperature simulation. There is no one way to include temperature, and this is going to be an important message.

And the reason is that extra energy goes into excitations of the system. OK? It goes into, you know, displacements of atoms, phonon modes. It can go into electronic excitations. It can go into defects. And you will have to explicitly choose which of these excitations you include in your modeling.

So every time you do a temperature modeling, most of the time, what that means is that you're going to have selected a series of excitations in which you allow energy to be pumped in for temperature, but that means you will have excluded other ones.

There is essentially no technique that allows for all possible excitations as you have in nature. So we will often go back between sort of the measure of temperature as temperature, or simply the measure of excitation energy as temperature.

Essentially, the amount of energy you're above the ground state can be used as a measure of temperature. And let me show that. Here's a little bit of statistical mechanics in a sort of classical picture. Some of you may have seen this already.

Some of you may have not. If you haven't seen this, just try to maybe just get the general idea. You know, in classical mechanics, you often write the Hamiltonian, which is essentially what's used to derive your equation of motion in generalized coordinates qi and generalized momenta pi.

And the reason you have both is, of course, because you have potential and kinetic energy. The momentum will represent the kinetic energy, and the generalized coordinates will represent the potential energy. So the reason they're generalized coordinates is because they don't just-- they're not just coordinates of atoms.

They could be linear combinations of those coordinates of atoms, for example. So the Hamiltonian is then essentially the potential energy as a function of the two and the kinetic energy as a function of the two.

In general, the kinetic energy can be a function of both, and the potential energy can be a function of both although normally you find that internal energy is only a function of the generalized coordinates. And the kinetic energy is only a function of the generalized momentum.

The equipartition theorem essentially tells you the definition of temperature. It tells you that the expectation value of, say, the momentum times the Hamiltonian derivative with that momentum, that the expectation value is kT and the same for the expectation value with the generalized coordinate.

And sort of you can kind of see a little bit what that means. The derivative of a Hamiltonian, which you can think of as an energy function-- the derivative of a Hamiltonian with respect to a coordinate, that's a force typically.

So this is a force times a coordinate, and here you have a force times the displacement. So these are energy terms. These are work and power terms. The expectation value is the time average expectation value.

This is the general statement not for quadratic Hamiltonians. You'll see we get a more useful and more common result. So a quadratic Hamilton is one where the Hamilton, the energy terms, are second powers of the coordinates. And this is what you very often get.

If you have a velocity, then your kinetic energy is mv squared over 2-- so the quadratic in the momentum the momentum being mv. If you had coordinates that move in a harmonic potential, a quadratic potential if two things were connected by a spring, an ideal harmonic spring, then the potential energy would be quadratic in the spring distance.

OK? So in that case, you can work through it because then the derivative of the Hamiltonian with respect to any of these things is linear in the coordinates. So since you multiply, you get something that's quadratic in the coordinates.

So then what you get is that essentially up to some constant the average of the squared coordinate, or squared momentum, is kT over 2. So what that's really telling you is that every degree of freedom you think of-- every qi, every pi is a degree of freedom.

Every degree of freedom contains a thermal energy equivalent to kT over 2. And this is sort of an interesting statement because it doesn't tell you anything about the stiffness of that degree of freedom. In principle, if I have even a stiff spring or a loose spring-- losing my words this morning-- have the same thermal energy at a given temperature.

Even though they will vibrate very differently, the stiff spring will not vibrate a lot, which will vibrate with small amplitude but at high frequency. OK? And the soft spring will vibrate at low frequency with lower amplitude. But so they all have the same degrees of freedom.

And so this is how you can calculate, like, typical heat capacity. So if you think of n atoms, how many degrees of freedom do you have if with n atoms you have three coordinates for atom? So you have three degrees of freedom for the coordinate-wise.

You have three velocity components. So you have 3n components from the velocity, so you have 6n degrees of freedom. So what's the thermal energy when you have 3n atom? It's 6 nkT over 2 of them. So it's 3nkT.

And that's the sort of classical result you find that the thermal energy is 3 nkT-- so what's the heat capacity? It's 3 times k. And those of you who've taken thermal from me, that's the law of Dulong-Petit if you remember that. It comes straight out of the equipartition theorem.

There's some subtleties. This statement is only true when the energy depends quadratically on these coordinates. So people often remember the statement. But now where it comes from?

So if your particle is moving in a non-harmonic, which you'll typically see if you start getting oscillations far from equilibrium-- around equilibrium, you tend to be quadratic in a potential if you go far away to see non-quadratic behavior. This will not be true anymore-- and the same thing for the velocities.

But there's maybe a more important subtlety. This is only true-- what we did now is classical mechanics, and the results we get are only true for degrees of freedom for which there's enough energy for them to be excited. And what do I mean with that?

We know from quantum statistical mechanics that the states are degrees of freedom, and the energy associated with them is quantized. So you cannot move from the ground state continuously in excitations. The first excitation from the ground state takes a finite amount of energy.

If you live on a temperature scale for which you don't have enough energy to get that first excitation, then the system will essentially not be excited in those degrees of freedom. And they will not contribute to the thermal energy. OK?

So if you live at, say, room temperature, that's 25 millielectron volts, and your first excitation is 1 electron volt, that degree of freedom is not going to contribute to the internal energy or the heat capacity. OK? So that's a bit of a subtlety that shows up with the equipartition theorem.

And this is if you remember why-- so the classic the classical heat capacity is 3 kb from what I've shown you before. But if you remember this is why if you plotted versus temperature, the heat capacity, it is 3 kb at high temperature.

But then it sort of drops off to low temperature because essentially the heat capacity coming from the atomic displacements gets frozen out because the atomic displacements are quantized. They're not continuous in a quantum mechanic sense.

So in some sense this is the subtext of the lecture. And if you're going to do a non-zero temperature simulation, I assume you do that to get a certain-- see a certain behavior or calculate a certain property.

You need to understand right away which excitations you're going to sample. Which excitations that are going to be present in your model to allow for the extra energy that comes with temperature?

So that means that you have to understand first the excitations that are relevant for the behavior you want to study, and you have to understand how they store energy. And I'll give you some examples of that. I'll try to make that a lot clearer in a second.

So, you know, how do materials store energy? Here's some possible excitation mechanisms. Electronic. You know, you're calculating quantum mechanics not only the atomic ground state-- you know, you relax all the atoms.

I think you've done that in the lab now to the lowest energy position. But you always, always find the electronic ground states, but at non-zero temperature you can have electronic excitations. You know, think of electron excitations across a bandgap as, if the bandgap's large, it's a small number of electrons.

But if you have a metal, you know, then the Fermi level is in the density of states. There's no gap there. So you can have, at any temperature, excitations of electrons across that gap. And I'll show you later some actual results to see how important or not important that can be.

Vibrational are the obvious ones. So that's the one you always think of. If you raise the temperature, what do atoms do? They move kind of. They start vibrating. That's kind of the typical obvious one. I know everybody's staring at this picture of springs going.

It's amazing how we can hypnotize you. And then there's of course more subtle excitations. If you have a compound, an ordered system where atoms are, say, ordered in a pattern, they might start inter-changing positions. So we think of these as configurational excitations.

The equivalent in polymers is that you change the confirmation of polymers at elevated temperature. So what kind of changes might you want to look at where you think you need temperature?

There's obviously things-- obvious one. Crystal structure changes. Sometimes you want to calculate what the structures of material, and it's already a really hard problem to do that at 0 Kelvin because it's a sort of non-local optimization problem.

If I give you gold, silver, and fluorine, and I tell you I'd throw these together in a bucket, what will they form? I don't even know they will form a solid, but let's say they did. That's a global optimization problem. You cannot solve that by local optimization minimizing forces on coordinates.

OK? So the zero-temperature problem is already hard enough. But then I tell you, well, the crystals structures we know can change as a function of temperature. So now not only do you have to find the ground state, but you have to figure out what are the sort of entropic factors, the excitations, that make free energy different from energy and therefore structures change as a function of temperature.

Chemistry can be very important. Changes in chemistry as a function of temperature-- this is sort of an often overlooked one. When you work with materials, many of them are open systems to the environment. A perfect example is an oxide if you take an oxide that's by definition open to the environment because it can exchange oxygen in the environment.

So most people who work with ceramics know that, if you heat them high enough, they start reducing. So they start losing oxygen. So this is one of these things where it's an obvious temperature effect. Properties. You know, there's obvious properties-- electrical conductivity, thermal conductivity.

All these things are very temperature dependent. They actually, in some cases, have trivial forms at low temperature. And then there's sort of trivial properties-- volume. Maybe you want to know what the thermal expansion does? So I'll go through several of these.

And the reason I do it is because there's no unique method to do all of these. So depending on what it is you're after, you're going to have to come up with different signs and a different toolset. So I thought I'd give you a very straightforward example.

Here's a work done I think by Andrew Huang-- I don't remember-- at Livermore showing the temperature dependence of the bulk modulus of copper. And what you see here is three different results. So that's the LDA result.

And you know what that is-- the local density approximation. The GGA result, notice that it's as you kind of expected and we discussed a little-- softer. And then the experiment is the green points. The experiment is right in between.

So let's say somebody asked you to calculate the temperature dependence of the bulk modulus of the material. How do you turn temperature on here? In essence, what I'm asking you is what is the essential physics that makes the bulk modulus temperature dependent?

Because if you know that, then that's the physics you want to include. OK? So you could say, well, you know, I'm going to do this brute force. I'll do a big molecular dynamics simulation. And after this we'll start molecular dynamics in a class.

And that's definitely a way that you would get the temperature dependence of the bulk modulus. But in some cases, you can do this much faster if you're sort of smart about what the physics is that going. Let me sort of show you what the question really is.

The bulk modulus is defined as the second derivative of the energy with respect to volume, but there's some subtleties there. Correctly, it's the second derivative of the free energy with respect to volume. Now you have to evaluate that derivative somewhere.

You evaluate it at the equilibrium lattice constant. OK? So what do you do when you do the bulk modulus at 0 temperature? What we do is you don't have the free energy. Well, the energy is equal to the free energy at zero-temperature.

So first of all, you take the second derivative of the energy not of the free energy. And you don't take it at the finite-temperature lattice parameter. You take it at the zero-temperature lattice parameter. OK? So we have two sources of error.

And one is they both require a different solution. We use energy to calculate the bulk modulus at zero-temperature rather than free energy. And we use the zero-temperature lattice parameter rather than the finite-temperature lattice parameter.

So the question is, which of the two is the dominant one? Or do they sort of-- if you're out of luck, they both kind of contribute the same effect, which means you need to deal with both of them. Well, that's what sort of this picture tells us.

What these two lines are is the bulk modulus artificially taken at different lattice constants. And the way you can do that is you can actually plot either energy or free energy as a function of lattice parameter, and then in the minimum, you have a bulk modulus.

It's the curvature there, but you could also get it, the curvature, at other lattice parameters. And so you'd get the bulk modulus as a function of the lattice parameter. And it's done in two ways. It's done by taking derivatives of the energy. That's the purple curve here.

And then the brown one is done by taking derivatives of the free energy. So what do you see? Well, what you see is that neglecting the entropy in the second derivative is not a big deal because, at a given lattice parameter, whether I take the derivative of the energy or the three energy is not a big error.

Where does the biggest error come from? It's that if I'm not at the right lattice parameter. So what causes the temperature dependence of the bulk modulus is the fact that the lattice parameter changes with temperature. OK?

And see now we've made a really big step forward, because if we're willing to work on the data approximation, then all we need to do is get the lattice parameter as a function of temperature and just calculate the curvature of the energy in that lattice parameter.

Now we have a good approximation-- the bulk modulus. Now this is a slightly recursive problem because, I'm going to show in a second that, if you want the lattice parameter, if you want to predict it, you'll actually need the entropy. So you need the free energy anyway.

But sometimes you have the lattice parameter. OK? Sometimes you just, you know, want the bulk modulus, say, at the experimental lattice parameter you've measured it. Now what I'm saying is, if you have that information, there is no need to do a temperature-dependent simulation.

You just take the curvature in that lattice parameter, and you're done. So the physics here that drives the temperature dependence of the bulk modulus is the thermal expansion. So here's sort of a zoology of numbers that I've taken out of this paper.

You know, I am a strong believer in numbers. That's why I feed you so many. I believe that science without numbers is out there. And so I'm sorry if that sort of bores you with numbers. But I do believe from actually looking at the facts, you separate fact from reality.

No-- fiction from reality. So here's sort of three different options for calculating the bulk modulus. So here the bulk modulus is calculated at the zero-temperature lattice parameter. That's what you would typically do.

Say, for a lithium, the lattice parameter that you calculate, this is in LDA, is 341. And you get a bulk modulus that's 15.4 gigapascal, and experimentally it's 11.6 just like you expect it's too stiff very typically in LDA-- same deal for the other metals although for aluminum it's not as bad.

OK. Then you could say, well, I remember what Sader taught me in the lecture. If I want to have the room temperature bulk modulus, all I need to do is calculate the bulk modulus at the room temperature by this constant. So that's what this is.

So you take the room temperature lattice constant. You calculate the curvature there. And notice that now you're on the other end of experiment. OK? And then you can do something else. You can do the full blown thing and actually try to find the equilibrium-room-temperature lattice parameter of your LDA, which isn't necessarily the real one. Remember?

LDA will tend to over-contract things and calculate the bulk modulus there, and that's what you get here. And, of course, this is better. And it's sort of typical because you're, in some sense, working in a consistent solution of the Hamiltonian.

You know, this is the minimum the Hamiltonian gives you, whereas at this lattice parameter, you're actually above the real minimum. And with this one, you're actually below it. So if you want to do the full solution, you would go and look at the thermal expansion to predict the lattice parameter.

Well, how do you do thermal expansion? Well, remember where thermal expansion comes from. Thermal expansion comes from essentially the volume dependence of the free energy. And that comes from two terms.

You have an energy that's volume dependent, and you have an entropy that's volume dependent. And the minimum volume is obtained by minimizing this with respect to the volume. There's another way you can show where thermal expansion comes from.

And here I have to brush up my own thermal again, but essentially you can show that it comes from the volume dependence that you can get it from only this term-- from the volume dependence of the entropy-- because the volume dependence of the entropy.

Here, you can use all this beautiful stuff of Maxwell relations and, et cetera, et cetera. You can write this as a chain derivative. And so what you get is that this is the compressibility times the dv/dt.

And so dv/dt, this is essentially the thermal expansion. And so what you see is, since dp/dv is always positive, pressure rises less negative pressure to get a smaller volume and increase the pressure. So the thermal expansion depends on how the entropy changes with volume.

So the critical part-- this is the hard part-- is calculating the entropy of the system and, not just the entropy, the entropy as a function of volume. And to calculate the entropy is exactly the same problem as dealing with temperature because the entropy is essentially counting the excitations your system goes into.

And temperature is the average energy of those excitations. But so to get entropy is the same problem as temperature. You have to know which excitations you have to deal with, and you have to have a model for them, and you have to equilibrate them. OK.

Where will most of that entropy come from? Typically, we think of thermal expansion largely caused by vibrational entropy-- by phonon entropy. So why does the entropy go up with volume, which is what gives you thermal expansion?

Because if you stretch the material, the bonds get slightly softer. So the vibrational frequencies gets smaller. They gets weaker. The bonds get softer, so you get more vibrations. Or in a quantized scheme, remember that the spacing between the zeroth level of a phonon and the first excited one is h nu, h times the frequency.

So if the frequency becomes smaller, your first excited state is closer. So you excite more. You have more disorders. You have more entropy.

And that's why most materials have positive thermal expansion, because essentially they get softer as you stretch them. So this is, for example, for lithium, the energy as a function of-- this is at room temperature-- as a function of volume.

So this would be the room-temperature-equilibrium lattice constant. And this is the free energy as a function of volume, so this would be the room-temperature-equilibrium lattice constant.

OK. We're not going to say much about the sort of subtle details about how you calculate thermal expansion. There is a sort of direct assimilation approach, and that will be the next lecture. So you could really just do molecular dynamics, which, again, we'll come to later what that exactly means, and equilibrate your cell size.

The other is to go more explicitly to a calculation of entropy as a function of volume. And so the way you would calculate the entropy is by calculating-- if you're only going to deal with vibrational entropy-- by calculating the phonon density of states.

So literally you would look at the quantized vibrations through the statistical mechanics on them. And those of you have taken 320 may have sort of remembered this and calculate the free energy associated with these phonons. And I'll say a little more about this later.

For example, this is a phonon spectrum for silicon. So you can actually quite accurately calculate it. Remember that to get the thermal expansion it's not enough to just phonons at one volume, because you get need to get the entropy as a function of volume.

So at a few different volumes-- that means you really have to do the phonons at several different volumes to get the thermal expansion right. This sort of looks hard, but phonon calculations have become quite routine. They're very well developed.

There's, in some cases, even, like, codes that do it almost semi automatic, which it's typically done either by a linear response technique where you sort of perturb a lattice really by a phonon with short wavelength, so that you can really calculate the energy of it.

So this is like setting up a supercell with a phonon in it is called a frozen phonon approximation. There are perturbation techniques, linear response techniques. And then some people do it in real space by calculating force constants and setting up force constant modulus for phonons.

So here's sort of the result that you would get in LDA and GGA for copper with phonons only. Again, the LDA number and the GGA number bracket the experiment, but I would say, overall, it seems to do very well.

Again, I think what you see is the LDA is a little on the stiff side. The GGA is a little on the soft side. It's very typical for these two in metals. So the LDA expands a little less. The GGA expands a little more.

But, you know, if you sort of look at the scale here, it's pretty amazing that this effect comes out quite well. But what about other things? Remember that the thermal expansion comes from how the entropy changes with volume, but that's the total entropy.

The calculation I showed you only had the vibrational entropy in it. OK? So any other level of excitations that changes with volume can affect your thermal expansion. And obvious one is the electronic excitation.

So any time you're not working on an insulator, you will start getting electronic excitations. And can they change with volume? And the answer is yes. What I'm showing you here is that-- and I must admit I forgot which material this is, which is really bad of me.

But I'm pretty sure it's the metal. So this is the ratio of the electronic thermal expansion, which is unfortunately called beta here not alpha, to the total extent of thermal expansion-- the electronic plus the vibrational one.

And so what you see is only at very-- and this is on a reduced temperature scale. This is the temperature in units of the Debye temperature. Think of the Debye temperature as sort of the temperature scale at which the vibrations become classical, which, for a lot of materials, it will be somewhere in the order of 200 to 500 Kelvin, so it's not 1 Kelvin. It's not a million Kelvin.

So what you see is only at really low temperature is the electronic thermal expansion important in these materials. Now why is it so important here? Anybody know? It's actually because the phonon one is really small.

Remember this is the ratio of the two. When you're at a real low temperature, the thermal expansion is very small in general. And the reason is you're essentially harmonic in your phonon degrees of freedom. So you have almost no thermal expansion.

If you have fully harmonic vibrations, then the entropy does not change with volume. Because if you are in a quadratic, your curvature is the same everywhere. OK? So your frequency is the same everywhere.

Your phonon frequency is the same, so your entropy is the same everywhere. So at very low temperature, you have almost no phonon thermal expansion. So the electronic one plays an important role, but all of them together are small anyway.

And here when you only at, say-- this is a logarithmic scale, so let's say at roughly half of the Debye temperature the electronic effect is almost meaningless. That doesn't mean there are not some esoteric materials where it does become important.

There are some materials-- sort of complex oxides typically where you can start inducing valence fluctuations at elevated temperature. But even 300, 400 Kelvin. And in oxides, if you change the valence of an ion, they sort of change size because the orbitals are very localized.

So there you can see a substantial effect of electrons on thermal expansion. But in many cases we wouldn't worry about it. So that was the example of thermal expansion. So there you have to worry about phonons, maybe electrons.

Here's another example where you would absolutely need a temperature-dependent study. Here's the phase diagram of titanium zirconium. Titanium and zirconium are both hexagonal close packed HCP at low temperature.

At high temperature, they're actually both BCC. And so this is the transition temperature on the titanium side. This is the transition temperature on the zirconium side. Let's say you wanted to know-- you wanted to be able to predict such transition temperatures.

So then you obviously need to get the subtle energy difference between HCP titanium and BCC titanium as a function of temperature. So you need to know what contributes to the entropy of these phases.

And you could say, well, you know, I'm really good at my initial stuff. So I'm going to do ab-initio molecular dynamics. And again I'm running a little bit ahead. Molecular dynamics is essentially you let the atoms move according to Newton's equation of motion.

So it's essentially doing the vibrations of the atoms. And if you do that, essentially, you include the vibrational entropy in the system, because you've included the displacement of atoms as your model for thermal excitation. So that's the same as vibrations.

So would you get that transition? I mean, these are questions you're going to be faced with as you apply modeling. It's when you deal with temperature that it becomes unfortunately rather important to understand the physics that drives your phenomena.

So you can sort make a blind choice. So let's say our ab-initio Hamiltonian was perfect, like LDA or GGA had no errors. And then we do molecular dynamics, so we let the atoms kind of move around. And we have infinite computer time.

So we can sample long enough and simulate long enough. Would we get that? I mean, the question really translates to, is that transition determined by atomic vibrations? Because that's the excitation mode you've allowed by doing dynamics of atoms.

I don't know. Do you know the answer? I don't. OK. I'm giving you a topic here. This is a hot topic of discussion among people who do alloy theory. What drives these HCP to BCC transitions, which are actually very common in metals or anything close-packed at low temperature to something less close-packed.

Most people will argue that it is the atomic vibrations-- that it's essentially a vibrational entropy difference. OK? The close-packed phase tends to be stiffer. The less closed phase tends to be softer. Remember softer means lower frequencies.

Lower frequencies means higher entropy. Higher entropy means you become favored at high temperature. I know it's a lot of steps to go through, but if you've never done thermal, you should get this by now. So if that's true, then our perfect ab-initio MD would get this transition.

OK? But what about other excitations? There are some people who claim that this is really driven by electronic excitations, because also BCC phases tend to have higher electronic entropy. And I'll show in a second why that is and how you could sort of know that.

If that's the case, then you shouldn't get this right just doing ab-initio MD. Then you should start explicitly including the electronic excitations, which can be done in ab-initio methods. But routinely it's not done because it's a pain in the butt.

And remember we often-- the problem is that, as we do modeling, we often tend to make approximations not because we can justify them but because we don't know how to do them differently. We tend to make approximations of convenience sometimes rather than justify them.

Here's another phase diagram. You think the same deal. Huh? This is copper gold. We just talked about titanium zirconium, and there it was either vibrations or electronic excitations. Here it's totally-- it's driven totally differently. These are FCC phases copper and gold.

And all these compounds in between are ordered FCC compounds, so they're really just ordering of copper and gold items on the FCC lattice. So this phase diagram is completely driven by what we call configurational entropy. OK?

So it's by the copper and the gold interchanging sites on the FCC lattice. So if you wanted to model this as a function of temperature, you would need a model that explicitly deals with that. And the vibrations, for example, may be much less important.

And actually they are much less important in setting the topology of that diagram. OK? So different problems, different physics. And it's going to lead to different models for dealing with temperature.

And here's the ultimate one. I thought that this didn't exist. For a long time, us alloy theorists have wiped away electronic entropy because we just don't know how to deal with it, and we don't like to deal with it.

So we write in paper saying that it's not important, but here's a system that's completely driven by electronic entropy. Now agreed it's slightly esoteric, it's vanadium oxide which undergoes metal-insulator transitions.

This is its phase diagram. And actually temperature pressure space since there's no compositional axis, it's a stoichiometric compound. And at low temperature and low pressure it's an antiferromagnetic insulator. If you actually raise the temperature, it becomes metallic and changes structure.

And it's completely driven by the fact that, in the metal, the entropy of the electrons is higher than in the antiferromagnetic insulator. So this is one of these things that, again, let's go back to our ab-initio MD. Like, if it was perfect, but all we did was let the atoms vibrate, we should not get this right.

OK? And it's amazing some people get it right, but, if you have completely the wrong physics that you allow in your model, there's no reason you should get the phase behavior right. This one I like very much because this is a topic that many people deal with.

This is a problem of interface energy. This is simply calculating the coherent interface between silver and aluminum. They're both FCC metals. So you put them together, and you calculate the interface energy.

This is something you very often need. And so this is that interface as a function of-- interfacial energy as a function of temperature. And the two types of dots are different interface orientations. And I forget. I think one is a 1, 1, 1 plane and one is a 1, 0, 0.

But I'm not absolutely sure, and I forgot to put the reference data. This is out of Mark [INAUDIBLE]'s work in Northwestern. But I'll put the reference on when I put the notes on the web. So obviously the interfacial energy goes down with temperature. OK?

And this is often something you need if you want to study wetting and all that kind of cool stuff and nucleation problems and stuff. So the question is, where does that temperature dependence come from? Because we've already had sort of three different types of things induced temperature phenomena.

So where does it come from here? Because if you don't know that, your modeling is going to be irrelevant. You can get a hint from the fact that it's getting really a lot smaller with temperature and actually kind of goes to 0.

What would happen to that interface as you heat it up? I think you know afraid to say. Diffusion. Yeah. It interdiffusers. Silver and aluminum are somewhat immiscible at low temperature. But when you heated up silver and aluminum interdiffused, they actually don't form a full solid solution until your way up, but they enter interdiffuse.

So essentially what's happening to that interface, that's sort of getting more diffuse. On the silver side, you have a lot more aluminum atoms. And on the aluminum side, you have a lot more silver atoms. So it starts to look less and less like an interface.

And actually by here near the interface you've reached almost perfect solubility, so that's why the interface energy is 0. So this is why-- it's because of the interdiffuse that the interface energy goes down. OK? So let's say that this was experimental data that you'd gotten.

If you had no idea about the physics from that interface, you would probably not apply the right simulation. Because if you took a hard silver aluminum interface-- you know, pure silver, pure aluminum-- and had all the beautiful stuff, you know.

Temperature dependent, had the vibrations, maybe even the electronic excitations-- you will not even get nearly close to the sensor, because the decrease in interface energy with temperature is purely driven by interdiffuse in this case. OK? Yes, sir?

AUDIENCE: [INAUDIBLE]? Or how is the [INAUDIBLE]?

PROFESSOR: Right. It's not just configurational entropy that's driving it up. Well, how is it defined? It's defined by looking at the energy of the two pieces together and then the energy of the two in the bulk with the same driving chemical potentials.

It's a good question because it is a little subtle, because they're both open systems. That makes it a little subtle. But what drives it is-- it's the configurational entropy that disorder it. But it's really the fact that-- it's really the temperature dependence of the interfacial energy.

Because they become so disordered, the energy of interaction becomes almost nil-- the effective energy of interaction-- because they basically look the same on both sides. You have effectively no interaction energy left. This is one that I used to ask about.

This is a thermal conductivity uranium oxide. This is a really old result. I have to find the reference. This is one of the very early simulation results. That's why it's on uranium oxide. I think I've told you this story. Potential modeling in oxide started in England with the nuclear industry, and they thought it'd be good to do this on nuclear materials.

So here's the calculated thermal conductivity as a function of temperature, and I'll tell you how to calculate this. To calculate this-- literally by doing a molecular dynamics simulation looking at essentially energy transport through phonons.

And the way you have-- thermal expansion is caused by phonons scattering. You have a phonon going through. And at some point, it needs to sort of-- that wave packet needs to become decoherent. It's scattered, so it can deliver its energy.

But essentially heat conduction comes-- at least some form of heat conduction comes from phonon transport. And that's what the calculated. What's the brown curve? And then the purple curve of the experiment-- and you see something kind of odd here.

There's obviously a whole region of temperature really good agreement. If I'd written a paper, I'd shown just that part. But then suddenly above 2,000, the two really diverge. And when you see something like this, it should always, like, make a light on there.

The physics must be changing. So uranium oxide is an insulator. The fluoride or antifluorite structure, I think. You're pretty high up here 2,000 degrees. What happens here is that you start getting an electronic contribution to the thermal conductivity.

When you start having hot electrons-- essentially electrons that can move-- they carry with them thermal energy. I mean, it's the basis of the thermoelectric effect. It's the Seebeck coefficient-- the Peltier coefficient. So at high temperature, you typically see a contribution of the electrons to the thermal conductivity.

In most materials at low temperature, the electronic thermal conductivity is small. It's largely dominated by phonons. But in a lot of materials it does start to show up at high temperature. And now what's high and low depends a little bit on whether you have a metal or an insulator.

If you have a really big bandgap insulator, they will never show up because you will never have mobile electrons. So that was kind of the intro to set the stage that you have to think of the excitations that you want to deal with when you work with temperature.

Now I was going to sort of quickly go through the two big lines of approach, and this is really a kind of table of contents for the rest of the course. One approach is to simulate dynamics. And this is what I've talked about before. This is what's called molecular dynamics.

And Professor Marzari is going to lecture on that for several lectures. Essentially, it's Newtonian motion for atoms. You say you've got a bunch of atoms together. On each of them, I can calculate the force coming from the other ones.

Force is minus the gradient of the energy. So it's the derivative of the energy with respect to the positions of the atoms. And then you say, you know, I know Newton's equation of motion. Force tells me the acceleration, so F equals ma.

So if I know the force, I know how much to accelerate that atom. And so you essentially keep track of velocities of atoms and positions, and you do a numerical integration of the Newtonian equations of motion. That's all molecular dynamics is.

And you do that a lot, a lot, a lot, a lot of times. So you track the dynamics of atoms. How do you set the temperature-- sort of a cool thing. I mean, Professor Marzari will go into this in much more detail. But the temperature is essentially the energy in your system.

But remember that Newtonian dynamics is energy conserving. So if you do your Newtonian dynamics perfect, then whatever energy you've put in in the beginning is whatever energy you'll have at the end.

So the temperature of your system is determined by the energy input you started. And what people typically do in simulations is that, if it's not quite right the way you started, you sort of adjust it on the go. You take out a little bit of energy. You put it back in if you want to adjust the temperature.

Again, you'll hear a lot more of that in the coming lectures. And there are flavors of this depending on how you get your forces. It's all the same molecular dynamics. But if you get your forces from empirical potential models, then it goes fast-- your force calculation-- because calculating the energy in empirical potential is fast.

If you get it from quantum mechanics, then it's called ab-initio molecular dynamics. And then it's slower. You have to get your forces from quantum mechanics, but it's the same molecular dynamics essentially. Oh.

So that's the first approach to simulate the dynamics. The other one, which we'll also deal, is to go through statistical mechanics. I mean, we sort of know how to average over microscopic degrees of freedom. And that's what's statistical mechanics tells us.

And in some cases that will be more advantageous. So you build some approximate model for the degrees of freedom. You try to get the energetics for those degrees of freedom. And then rather than simulating, you sort of do some more analytical or numerical integration of them and get, say, temperature dependent properties.

And we'll show you that as well-- how to do that. So you can kind of see that this is what we're going to cover. I just wanted to give you perspective on what's coming. The sort of first piece here, this is a direct simulation. We do molecular dynamics.

And then the rest here, kind of largely this piece, is sort of more, I would say, stat mech and making relevant models for your degrees of freedom that are not obtainable by direct simulation. And then this, the rest here, is kind of a bunch of case studies and separate applications.

OK. I wanted to get out of this and show you a movie. So here's a molecular dynamic simulation of a planetary gear. I don't know if you've ever seen this one. It's one of the totally crazy ones.

So somebody had the idea that you could make-- in the world of nano, you could make a nano gear. So these are actually molecules-- let me run it again. You see the molecules in the center turn faster, and they kind of rotate over the other ones.

And the outer circle runs slower. And this is done by solving Newton's equations of motion essentially for these molecules. I'll show you one more. No, it's not the right one. Oh, here we go. This is buckyballs shot onto a silicon surface.

So you can sort of see they kind of-- some stick. Some get shot at it. Some stick, and some bounce back up. And they kind of vibrate on the silicon surface. And so this is all done simply by integrating Newton's equation of motion. From

OK. Let me go back to the lecture before you get totally enamored by the movies. OK. I'm going to say a little bit about sort of which route to go when you want to deal with temperature.

You know, an obvious way is always to do the direct dynamics through molecular dynamics. It's a powerful approach. It's, in particular, powerful to see phenomena that you didn't even know were there, because in almost all the other approximate approaches we're going to do for dynamics or temperature, we're going to assume that we know and understand the phenomena that are responsible for the macroscopic behavior we see.

In some sense, molecular dynamics is the only way where you don't presume much. I mean, you presume something about your Hamiltonian, your energy function. But otherwise, you just kind of let the system run, and you see what happens.

What you have to keep in mind, of course, is that the timescale over which you run is extremely short. I mean, you're now living on the time in the world of atoms. And atoms vibrate, let's say, somewhere between 1 terahertz maybe 10 terahertz.

That's actually pretty high-- 10 terahertz. So that would be 10 to the 13 times per second. And remember you have to integrate equations of motion. So think of an atom vibrating in some well or against another atom.

You probably need at least 10 discretization steps along that one vibration. So your time step is of the order of 10 to the minus 15 seconds-- 10 to the 14 seconds. That's sort of the order of magnitude. For different materials and different algorithms, it will be different.

That's 10 femtoseconds. So you have to kind of do an iteration step every 10 femtoseconds. So that means even if you do 100,000 time steps, you still only have 1 nanosecond of time. OK? And so you'll see typical-- and these simulations are in the order of picoseconds to 100 nanoseconds.

Those are really long simulations. And it's amazing, but you can see a lot of important physics recurring on that timescale-- sometimes the reaction between molecules or absorption resorption phenomena.

If you sort of want to do, and the first thing you should do is a back-of-the-envelope calculation is the phenomenon that I want to see, or that I hope to see, does it live anywhere near this time scale? And sometimes you can give the system help.

You can raise the temperature, because stuff will happen faster. But you should realize that there's certain stuff that you do not see on this time scale at all. There are people who study solidification out of liquids within d. And, you know, there may be even some here.

It's stupid period. You do not see anything relevant on this time scale that has to do with any real solidification. The only thing you can do is enormously super cool a liquid, overdrive it, and really sort of induce extremely fast crystallization which has nothing to do with real solidification.

And so the one example I give is diffusion in solids. You can do dynamics on fast diffusers, but most things are not fast diffusers. So I did a-- I show you the kind of back-of-the-envelope calculation you should do.

A diffusivity is a length scale squared times the jump rate, where a is essentially the distance which you jump. OK? So if you take a relatively fast diffusion at room temperature 10 to the minus 14 centimeters squared per second, it's a respectable diffusivity at room temperature.

Most materials, you would never see that. It's not quite a fast diffuser yet. This number would be more at 10 to the minus 11, 10 to the minus 12. And you take a jump length of 1 Angstrom. So that's 10 miles 18 centimeter. Then you can show that the jump rate is about 100 per second.

So if you have a 10 to the minus 14 diffusivity, an atom jumps about 100 times per second, or about 1 times per 10 milliseconds. So do you see now the enormous disparity in time scale from 1 nanosecond to 10 milliseconds?

So you could wait forever in your dynamics to see things if you're sort of not smart. You'll see similar phenomena. Like, you could study. You could use dynamics to look at, say, surface deposition. Let's say molecular beam epitaxy or something like that where you lay down layer by later.

If you actually do the math in a real MB experiment, you tend to lay down about one layer per second. That's the order of magnitude. So you can actually never get that by direct simulation. So what do people do?

Well, people do the short simulations to extract useful information that they can then use in other models to get the real time scales that they want. And so that's some of the things we're going to teach you. OK. Probably going to run out of time here.

But the other option-- so if you don't. If you cannot do the dynamics-- is to use statistical mechanics and sort of integrate as far of the length scale and time scale as you can.

And just those of you who may have never seen statistical mechanics, I'm going to now totally confuse you. But the essential ideas are sort of in these four equations that you now consider a macroscopic state as a collection of microscopic states that you can fluctuate in between.

The probability for each of these microscopic states is proportional to the exponential of minus beta that the energy, and beta is 1 over kT. So what does that mean? Microscopic states with low energy show up with high probability, and microscopic stage with high energy shell up with with very low probability.

Makes sense. The sum of all these unnormalized probability terms is called the partition function. And you can show that knowledge of the partition function is exactly the same as knowledge of the free energy. The free energy is essentially the logarithm of the partition function.

And the entropy is the product of the probability times the log of the probability summed over these things. And again, that makes sense because if you are in one microscopic state, then pi is 1 for that micro state and 0 for all the other ones.

So 1 log 1 0. And 0 log 0 is also 0. So your entropy is 0. So you have no disorder. So if your macroscopic state leads to only one microscopic realization that you can beam, then you have zero entropy. It's when you can be in multiple different microscopic states that you have entropy.

So you can think of atomic vibrations, for example, as really going between different microscopic states of the positions of the atoms. You can think of electronic entropy as going between different electronic microscopic states.

OK. So I want to show you one example and probably end with that. Electronic entropy. It's one that, on the face of it, looks very difficult to deal with, because you have to deal with the excitations of the electron. But actually in many cases it's much easier than you think.

If you solve the Kohn-Sham equations, you get the eigenstates i and the eigenvalues for the energy. And typically we tend to treat these as one electron eigenstates. So we sort of fill-- we think of filling these with an electron, where we have a bunch of them, and we fill them up to the Fermi level.

There's actually nothing in the theory of density function that tells you that they are one-electron eigenstates where we treat them like that. So if you think of-- I have a bunch of. Let's say I draw them as discrete levels, a bunch of energy levels-- E1, E2.

And, of course, in a solid, this will become a density of states. What's the level of uncertainty? What do I need to know to characterize my electronic configuration? All I need to know is, is the level occupied or not?

So I can characterize my electronic state by saying which levels are occupied or not. So if you do this on a solid now where you have a real density of states-- and I've shown you one here. So this is a density of electron states. This is energy.

And forget the bottom piece here. Just think of the top. These are actually spin polarized. One is in the up dos, and one is the down dos. So my Fermi level-- I don't know. It's actually at 0 here, but that makes it less interesting. Let me assume the Fermi level is at minus 2 here.

So at 0 temperature I feel all the states below the Fermi level. OK? At elevated temperature, you know you apply the Fermi-Dirac distribution function, and so you'll have a filling that's maybe a little different.

And I'm kind of exaggerating. The width of the Fermi-Dirac distribution function is not nearly-- so it's 1 here and 0 here. It's not nearly as wide at reasonable temperature. So how can you get the entropy? Well, you can get the entropy from the probability that each electronic state is filled or unfilled.

OK? So if f is the probability that it's filled-- so this is the entropy coming from when the state is filled, the entropy coming from when the state is unfiled, and the sum over all these possibilities is the electronic entropy.

So if, in ab-initio calculation, you calculate the density of states, you actually have a trivial way of getting out the electronic entropy. And this is-- getting the density of states is really rather simple.

So you just smear it with the Fermi-Dirac distribution function, and you sum-- or you do this as an integral typically rather than summing. And you have the electronic entropy. It's very simple to do.

So just as sort of a caution, what I should tell you is that usually this works very well, but the fundamental underlying approximation lies in the starting point. It's in assuming that the epsilon I's, or the eigenvalues, are the eigenstates of the real electron system.

So what I'm essentially saying is that if you excite an electron or remove one from the occupied states, you're essentially perturbing the system exactly by that energy of that state that I took out. If I have a Kohn-Sham eigenvalue of 3 electron volts with some eigenstate, and I put the electron in there, I'm assuming that that's exactly the perturbation that a real electron will do in the system.

If you go back at the theory of density function, there is nothing in the theory that tells you that that's right. OK? Remember density function theory is based on the density? And the Kohn-Sham idea was essentially the smart way of writing the density in terms of things that look like one electron orbital but they aren't.

OK? So an eigenstate densely function theory is not a one-electron orbit. We just hope it is, and it sometimes looks like it is. But it's not. And if it's not, then you can describe excitations of electrons as going between one Kohn-Sham eigenstate to another Kohn-Sham eigenstate because they're not one electron eigenstate.

OK? So what we are doing is completely hokey. It's just that, in reasonable materials and metals and semiconductors, they don't look very much like one-election eigenstate, and this stuff works. Where it breaks down is typically in materials with very high degrees of correlation, transition metal oxides, where these things don't look at all like one-electron eigenstates.

Because what happens in those is, as soon as you do an electron excitation, the whole underlying eigenvalue spectrum changes. OK? And remember then it's not an eigenstate. An eigenstate is supposed to be orthogonal to the rest.

So if filling or unfilling it changes the occupation of all the other ones and the wave function of all the other ones, then it wasn't an eigenstate. So there this stuff fails. I'm kind of out of time. So what I was going to do was show you the next picture, which you may just have a look at.

It shows just the contribution of the electronic entropy and the vibration entropy in metals as a function of temperature, so you can sort of see how they compare to each other. So actually let me end here. And we'll see you back on Tuesday.