

GERBRAND

What we're going to do is work through several case studies so you'll see it get more and more applied. I'll start with doing a little more of alloy theory and cluster expansions, which is really a follow up on the Monte Carlo work we did and you'll see how that fits in with course grading. It's essentially the part we didn't get to when I finished. And then I'll do a variety of things on Tuesday.

CEDER:

One of them is [INAUDIBLE] accelerated molecular dynamics technique, which is sort of a specialty topic but it gives you an idea of where people are trying to push the boundaries. And then on Thursday, we'll actually have a guest lecture. We'll have Chris Wolverton from Ford Motor Company lecture. It's kind of cool, he actually lectures from Ford so you're going to see him project on the screen. He does that every year and people usually find that a really good lecture.

He essentially talks about how they use modeling in general but in particular, first principles methodology at Ford Motor Company. He has a few really cool stories. He won't tell you how much money they save the company, but it's actually in the tens of millions of dollars so that's why they're still there. He won't tell you that, but I can tell you that before his talk show. That's on Thursday, and then we go into the last week and I think we have only one more lecture left there.

So one of the things I want to do today is be a little more systematic about some of the models we've introduced for your lab. We introduced this lattice model, for example, to model hydrogen absorption on palladium and several people have questions, you know? How do you know that this is a good model?

Those lattice models actually occur all the time in a solid state structure equilibration thermodynamics and I'll go through them a little more systematic today and then towards the end prove to you that they actually are a very accurate representation of the thermodynamics of real systems. So let me show you again where they will tend to show up.

I introduced the lattice model too simply as a magnetic spin model, and that's how we started. You model a magnet and you have up and down spins. But as soon as you realize that those spins are really just general occupation variables, they tell you the result of a binary outcome, you can also use them to represent a whole lot of other things.

For example, if you look at superstructures in the copper-gold system-- copper is FCC, gold is FCC, all the structures in between are in an FCC lattice. So if you really want to know how they look, the question is, how do you distribute copper and gold on an FCC lattice?

So really all you need to know is on every lattice site, is there a copper there or is there a gold? Again, that's a binary outcome, so you can represent that with spin variables. So essentially the energetics and thermodynamics of this will be represented by a spin model, also by a lattice model. It doesn't necessarily have to be atoms, it can be atoms and vacancies.

This is an example I'll come back to later, this is lithium cobalt oxide. The purple things up there are lithium, red is oxygen, and the green is cobalt. You can actually extract lithium from these materials, and so start removing some of the lithiums. And so essentially what you have then is a distribution problem of lithiums and vacancies on a given set of sites.

In this case, those planes that you're looking head on are actually triangular lattices, OK? And I'll come back to that as well. So here the spin variable would indicate whether the site is occupied by a lithium or not. You can do surfaces. Lattice models are classics for surfaces. Here's a typical 111 surface with stuff absorbed on it.

So if you look at an FCC metal like platinum and you look at the 111 surface, the absorption sites have the same symmetry. For example, the hollows, which I've indicated here. So the hollows form a triangular lattice, OK? And so if you look at absorption, again, all you need to know is, is there something at the absorption site or not? So again, you would use a binary variable.

So what I first want to show you is that an Ising-like model or a lattice Hamiltonian, as you've seen before, is actually a very natural way of writing the energy and actually, I extend that later to any property of configuration. So I'm going to do that a little more formally than we've done before. Essentially if you want to study all of these models that are lattice models, you need some Hamiltonian for them.

And so I'm going to use this sigma between these squiggly brackets as a variable to indicate the configuration on the whole lattice. So the variable sigma is essentially a combination of sigma 1, sigma 2, sigma 3. So you can think of as a vector in the configurational space of all the occupation variables. So what you need is a practical way of writing that configurational Hamiltonian, and the way you make you write almost any practical Hamiltonian or any practical energy function that you want as accurate as possible is that you write it as some expansion in a suitable basis set of functions. OK?

Think of how you explain wave functions in plane waves, how you sometimes expand in sines or cosine. The typical basis that's used here is what's called cluster functions. A cluster function, which I'm going to label by phi alpha, is essentially the spin product of all the sites of that cluster alpha.

So if I make a pair cluster, let's say this one, i and j, the cluster function phi alpha then is sigma i times sigma j. But I can make many other ones. I could make a four body one, they don't even have to be connected. And so for the red, phi alpha would be-- let's call it i, j, k, l, would be sigma i, sigma j, sigma k, sigma l. OK? I think the computer gets more tired as well as we proceed. OK.

The first thing, so when you expand in a basis, you want to make sure that the basis is complete to describe your full space and then sometimes it's useful that your basis set is orthogonal, although that's not really necessary, but it can be practical. It's easy to see here that it's complete. I mean, if you have n lattice sites, what's the size of the space? What's the size of the configurational space? Well, it's 2 to the n.

On each lattice point you can put two things, essentially, so the number of combinations is 2 to the n. Well, how many basis functions are there? Well, if you think of what's the choice of basis function, it's essentially picking a combination of sites. So that's also 2 to the n. There are 2 to the n combinations of sites because it's also a binary problem. For every site, you include it or you don't include it in the basis function.

So there are 2 to the n basis functions. I'm going to show you in a second that they're orthogonal. The basis that is complete, in a very trivial sense, there are as many basis functions as states. OK? So you may ask yourself, what have you won? We're actually going to expand things in this basis set, but there's as many basic functions as state.

Well, the hope is that we can formulate these basis functions so that whatever quantity we expend is more rapidly convergent. So let me prove you that the basis set is orthogonal if you take a suitable scalar product. The way I'm going to define the scalar product between, say, any two functions of configuration, which I'll call f and g . OK?

I'm going to define the scalar product simply as the sum over the configuration space properly normalized. OK? So this is equivalent to if you're working in a continuous space defining your scalar product as the integral of the product of the two functions over the configuration space, and you might be more familiar because we have a discrete configuration space. That integral is essentially a sum over all the possible configurations.

Let me prove you that the functions are orthogonal. First you can note that for any basis function, if you sum its value over the space it gives you 0. And you can easily see that-- let's say that function would look something like $\sigma_i, \sigma_j, \sigma_k$, maybe even more terms, essentially these terms will alternate as having plus and minus 1 and so they will sum to 0. Using that, you can easily prove that the functions are orthogonal.

So if you adopt the two basic functions into each other, so you're essentially summing their product, OK? So let's say we have a cluster α and a cluster β , maybe this is the cluster α and maybe this is the cluster β . OK, so they have a partial overlap. So you can write the product of α and β then as first as the product of the spins that are in both α and β .

And for those, you get σ_i^2 , and then the spins that are only in α and the spins that are only in β . Of course, σ_i^2 is always one because σ_i is plus or minus 1 your occupation variable in a spin model, so essentially this goes away. OK? And this is simply a function. You could write this as the basis function of $\alpha \cap \beta$. OK?

You're essentially here multiplying all the spins that are only in α or only in β , but not in both. So that's trivially 0. So the only time that this product is not 0 is when this part doesn't exist, so that means when all things are in both α and β . OK? So this product is equal to the Kronecker delta. It's actually not this product, I'm sorry.

It's this product summed over all configuration, so I should correct this. It's $1/2^n$ and then that should go in front of here. So the basis set is orthogonal, OK? So we have a very practical expansion now, we have a complete and orthogonal expansion.

So one of the first things we're going to do is expand the energy in this basis set, because the energy is a quantity we'll need in things like Monte Carlo simulation. But you should right away see that this is much more powerful, huh? You can expand any property of configuration in this basis set. They may have different rates of convergence, but essentially this is a basis set in which you can expand anything that depends on how you arrange the items on a lattice.

So people have done things like expand the band gap of materials. And if that converges fast, you have a model that you can rapidly search then. People have expanded things like the vibrational entropy in it, people have expanded things like-- even tried to expand things like connectivity, and so which goes not as well. But essentially anything that becomes a property of the arrangement of atoms on that lattice, so in that substitution space can be expanded.

For now we're going to stick to the energy. So if we expand the energy as a function of configuration, so we get an expansion over all the possible clusters α , the basis function, and then the expansion coefficient, which we'll call v_α , and so of course the expansion coefficient that's going to contain all the energetic information. And I'll show you in a second how it's defined, but first I want to write out what this cluster expansion, as it's called, looks like.

I mean, if you pull out what these cluster functions are you essentially get something that looks like this. There'll be an empty cluster function, so that's this term. When α is just a point you'll have this term. So this is almost like point site energies or chemical potential like. If α s are pairs, you'll get the standard pair term that we see in lattice models, but this keeps on going. You'll have triple terms, quadruplet terms, and so on.

So essentially the cluster expansion is the same as what's called a generalized Ising model or "Easing" model, OK? So where you go away from relying only on two body terms to higher body terms as well, and what it essentially is is a generalization of that approach. I'm sorry, I'm going back and forth.

I'm interchanging and using e and h , h for the Hamiltonian, e for the energy, and one of the reasons is that I cut-and-paste equations from various PowerPoint slides and wasn't quite able to modify them due to the great Microsoft software. OK. But I think what's the really nice thing is that this formal expansion will allow you to define exactly what these v 's are because you may have asked yourself when you were doing in the lab and are probably still doing the hydrogen on palladium, what those interactions actually are?

Are they sort of physical potentials between the hydrogen or not? You'll see in a second that they're not, they're much more complex than that. OK.

So whenever you write an expansion in orthogonal basis functions, how do you get the expansion coefficient? Well, the expansion coefficient is essentially the projection of your function onto the corresponding basis function, and that's what you'll see here and you'll get something really interesting come out. So the way to know how a v_α is defined, you can essentially take the scalar product of this expression on both sides with ϕ_β .

So you multiply with ϕ_β and then sum over all configurations. So you're essentially taking the scalar product here, ϕ_β, ϕ_β . On this n , we know that this is the Kronecker delta, so this is v_β . OK? So we've projected out v_β out of this expansion, and then we have to work out essentially what this is to know exactly what v_β means. And you'll see this is the projection of the energy onto the basis function [ϕ_α] but it's interesting when you work that out what it physically looks like.

So that's what I'm going to do here, so I've written out the scalar product here. OK. What I'm going to do is I'm going to-- this is the sum over all possible configurations of the lattice. Since I have a basis function ϕ_β here, I'm going to separate that into a sum over the configurations on the spins that belong to β . So let's say this is β , sum triplet. So I'm going to sum over those spins and then separate that and sum over all the other spins on the rest of the lattice.

So I've written that here as a sum over all the i 's-- sorry, this should be \sum_i that are in β and summing over all the configurations, you get that by just making all the σ_i plus or minus 1, OK? And then summing over the configuration space minus what lives on β , and this is a bit shorthand notation. It's a little hard to write formally what you're actually doing here. But those of you who are familiar with quantum mechanics and so may recognize this as a reduced trace.

You're essentially tracing over all of the configuration space except what lives on β , OK? OK, so what I want to do is work out the example for pair interaction since the math is about just the right level to do in class. The three body gets a little harder. So we're going to work this out when β consists of sites i and j , so it's a pair interaction.

So you see that in this first term I'm going to get four terms because I have two sides. I can do a or b , let's say. Let me call plus 1 is $[\sigma_i = +1]$ and minus 1 is $[\sigma_i = -1]$. So let's say I do the term $\sigma_i = +1, \sigma_j = +1$. OK? So the basis function is going to be 1 because I have plus 1 times plus 1, and then what is this really? Think about this.

So when i and j are plus 1, that means I essentially decorate this cluster of sites i and j with an a atom and an a atom. So it's an aa pair. But what is this term? Essentially I'm summing over all the configurations outside of this cluster, all the spins. So this is essentially the average energy of the configuration with aa pairs at i, j , because look at what I'm summing over.

So I'm summing it over the energy of all the possible configurations, but it's configurations that have a fixed β occupancy because I'm not summing over β . So this first term is essentially what I'll call EAA averaged. So you see what it is? It's I put a at i , a at j , and then I essentially trace over the remaining configuration space. So I put all possible configurations on the rest of the system and calculate that average energy. OK?

So you can do that for the other terms. So I'll write out explicitly now V_{ij} , because β is a cluster consisting of i and j . We're going to find this n_β is the number of sites in β , so we get $1/4$ because the number of sites is 2. So the first term I got was EAA. OK. If I do $\sigma_i = -1$ and $\sigma_j = -1$ I'll get, again, plus $4 \phi_\beta$. So I get plus EBB.

And then if I do i and j a different spin, so plus 1 and minus 1, I'm going to get minus EAB minus EBA. OK? Does everybody roughly see where that comes from? So I essentially explicitly evaluate the four terms in that sum and every time, I essentially get the average energy of that configuration.

If your system is symmetric, if i and j are equivalent, then these are the same. OK? If you have translational symmetry between i and j , so then you'll get something like minus $2EAB$. And those of you who've seen regular solution thermodynamics may recognize this very much as a sort of effective pre-factor in the entropy of mixing. You'll get something like the AA interaction plus the BB interaction minus 2 times the AB interaction, just sort of pointing that out. But it's sort of interesting, this is the formal definition of the expansion coefficients.

First of all, note that they are not potentials, they are energies. OK? And they are energies of configurations. So this is not the potential between the a atom at i and the a atom at j , OK? The reason is that the configuration around it modifies that interaction substantially, OK? Typically what you'll find is that these v 's are much, much smaller than direct interaction energies, and the reason is they contain enormous amount of cancellation.

If you think about it, this is AA energy, BB energy, and minus 2 times the AB energy. So they're often several orders of magnitude smaller than direct interactions if such a thing were to exist. The other thing is that indirect effects can be included, like if you truncate your lattice model at pair terms. OK? Because it's an infinite expansion, at some point you'll have to truncate it.

You are still including many body quantum mechanical effects because the many body quantum mechanical effects enter in the calculation of the energy of these, OK? So what you should understand is that the level of physics or the level at which you truncate the cluster expansion, the lattice Hamiltonian, has no correspondence to the level of physics for, say, your electron physics. OK?

Let me give you an example just to think through what these things are. Let's say you had a model on a lattice where things just interacted with pair potentials, OK? So a very simplistic energetic scheme.

If the atoms don't relax away from the lattice points, then these energies can really be written as a pairwise sum because then the total energy of a system is really the pairwise sum between all the pairs. Then now what you're going to find is that essentially, all of the pairwise sum cancels in this difference except the AA interaction, the BB interaction, and the AB interaction.

So if you do just a pairwise energy expression-- and should try this one, it's very pedagogical-- and the atoms don't relax away from the lattice, then you'll find exactly that these are pair energies, pair potentials. Now as soon as you let the atoms relax away from the lattice point, like maybe the a one is bigger than the b one, even in a pairwise energy Hamiltonian these will not be the pair interactions because you cannot remove from these energies.

You cannot cancel the energetic effects between a and the rest of the material and b and the rest of the material, OK? So the really cool thing about these expansion coefficients is that they will include all that extra physics, like relaxation away from the equilibrium lattice site, and I'll come back to that later. Come on. Shouldn't write so much. I think class is over, no? Ah, there we go.

So I'll go back and forth between writing the cluster expansion in terms of the ϕ_α , so the basis function, or being more explicit and writing the basis functions out as the products of spins right away. So all your physics of your system, all the energetics sits in those coefficients v of course. So the question is, how do you get them? Well, there's a lot of different ways that people have tried to get them and it depends a little bit on the problem.

For simple Hamiltonians, you can often really just calculate them directly by their definition, essentially as the projection of the energy onto that basis function. You can do that for pairwise Hamiltonians even including forced relaxation. So even when atoms don't sit on a lattice, you can do it by Fourier transforming, and that's in one of the references that I give you. But as soon as your physics is very many body, it becomes very hard to do it analytically.

There are people who do it numerically very much according to the definition. So you literally take a system and put an AA atom here and then calculate a lot of different configurations around that. And so literally start to approximate those averages, EAA average, EBB average. The practical way that people pretty much do it now most of the time, which is also the least elegant way, is that essentially, you fit your whole expansion at once to a large number of energy calculations.

So the first thing you have to do, of course, is to truncate your Hamiltonian because this is an infinite expansion. You should really see this, this is a trivial expansion. You're expanding 2 to the n quantities and 2 to the n basis function, OK? So it's only practical if you truncate it, but once you truncate you have a finite set of coefficients to determine. So what you can do is literally calculate the Hamiltonian, or in this case, the energy for a lot of these configurations.

Let's say there's 20 of these v 's you have to determine, then essentially if you have the energy at 20 different configurations you could determine the interaction coefficients. Typically people use very over-determined sets. So if you want 20 of these, you may do 50 of 50 energies and do a squares fitting just to make the determination of interactions much more stable. So that's essentially how it's done.

You have an energy method. I wrote it as a first principles method here, could be any energy method. You calculate a lot of energies of different configurations and then you fit the Hamiltonian and you plot the interactions out. It's really that simple. And once you have your lattice model Hamiltonian, you now know that it's fairly easy to equilibrate that.

So you can get ground states out of that, you can do a Monte Carlo simulation on it just like you're doing in the lab, and then get phase diagrams thermodynamic properties out. So essentially you can get both the structures and all the thermodynamic information and more important, as I'll show you in a second, all the information you can derive from the thermodynamic information, which is something people often forget. So what I was going to do is show you a few examples of how this is used and then in the end, essentially come back to some of the theory of that.

So here's one that shows how elegant it can be. What I'm going to show you is a calculation to try to get the solid part of the calcium oxide, magnesium oxide phase diagram. This is perfect for a cluster expansion. Calcium oxide and magnesium oxide are both rock salt structures.

So a rock salt is essentially an FCC lattice of oxygens, which I think are in this picture, the purple ones. And then the cations, either calcium or magnesium, sit in the octahedral sites, octahedral interstitials. And the octahedral interstitials of an FCC lattice also form an FCC lattice. So you can think of rock salts as two interpenetrating FCC lattices, one of cations, one of anions.

So when you mix the two, you're essentially asking the question, at each interstitial is there a calcium or a magnesium sitting there? So again, you have a perfect binary problem. It's a lattice model. But this already illustrates a nice thing of this approach, in your whole lattice model, the oxygens won't appear.

So the variables you use are only going to be defined on the cation lattices. So the cations are the smaller ones. So on this one, you'll have σ_i say plus 1 and on this one, you'll have σ_j equals minus 1. So you're essentially going to express the energy solely in terms of spin variables on the cation sites, and that makes sense. There's nothing unknown about the anion sites.

They're always fully occupied, OK? So the anions will show up when you calculate the energies of these configurations, of course, but there's really no need to expend the energy in their occupation variables because they're always there. They're always minus 1. So you see, we're going to end up with an effective Hamiltonian between the cations that includes the energetic intermediary of the anions without explicitly including them.

So how do you do something like that? Well, you truncate your Hamiltonian [? number, ?] it's always going to be the same procedure. You calculate the energy, a lot of different configurations. This is old work, so it wasn't fully quantum mechanical. There are two sets of points-- so this is actually energy-- I'm sorry, this is in electron volt per cation. The blue points are done with empirical potentials.

Remember, we did that all the way in the beginning of class? These are essentially Buckingham potentials with electrostatics and shell models attached to them. And the open red circles is a self-consistent PIB, which is a really poor man's quantum mechanics. It's a sort of intermediate between Thomas Fermi theory and better quantum mechanics, but think of it as a poor man's quantum mechanics, and these are mixing energies.

So they're referenced to 0 here and here, that's why they have that parabolic shape because you force them to 0 at 0 and 1. First of all, they're positive so that means the two things don't want to mix, which is in accordance to the phase diagram, but you may notice something already that's kind of interesting. The quantum mechanical values actually tend to be lower than the potential ones, and this is very common.

Typically a better energy model will lower the energy of low symmetry states more than a bad energy model. And the reason is that usually a good energy model has more degrees of freedom, OK? If you think of your potential model has no electrons in it. There can be no electronic relaxation by definition in a model with empirical potentials because you don't have electrons.

A quantum mechanical model has electrons in it so you have more ways for the system to relax. And the lower the symmetry, the more a system can actually use those methods. If your system is fully constrained by symmetry, then there's barely any way the system can relax. OK? So this is something you very typically see.

If you have energy methods that are fitted to perfect bulk properties, so they're fitted in high symmetry environments, if you then use them on defect properties, the better the physics of the model or the more degrees of freedom usually the lower the defect energy will be. But anyway, so you use this. You'd fit your Hamiltonians to it, your cluster expansion to it, and you'd get a bunch of pair interactions.

So the first six are pair interactions. Again, blue-- oh, dear. Changed the color scheme here. Blue is the quantum mechanics, now the open circles are the potentials, and then there are some triplets interactions and some quadruplets interactions. And notice that the first pair interaction is negative, so what does that mean? That means the nearest neighbor's spin product wants to be positive, get your energy down.

So that means like things want to be together. Plus ones want to be with plus ones, minus ones want to be with minus ones, so that's phase separation. You can already see it from the interaction, but notice the scale of the interaction. So this is now about 40 milli electron volt per cation. So this is a far cry. This is much lower than any direct electrostatic interaction or even any direct evaluation of a Buckingham potential because you have effective interactions.

Actually, if you think of the definition of the effective pair interaction in this case-- so let me write it out. So V_{ij} would be $1/4$. Let me actually put calcium and magnesium on here. So it would be $E_{\text{magnesium}}$, magnesium averaged plus E_{calcium} , calcium minus $2e_{\text{calcium}}$, magnesium on the i, j sites.

So you should think about where that effective interaction comes from because look at the electrostatics. Does the electrostatics contribute? Magnesium and calcium are both 2 plus. Essentially if you froze those two on their lattice sites, you'd have in the first term 2 plus 2 plus plus the 2 plus 2 plus minus 2 times the 2 plus 2 plus. So the electrostatics completely vanishes from the effective cluster interaction.

So it's interesting, the fact that you do this projection essentially filters out the largest part of the energy here in the system, but because it's the same for both magnesium, magnesium and calcium, calcium and magnesium, calcium interactions it's irrelevant for the configurational thermodynamics and the projection essentially filters that out. So where does the remaining V_{ij} come from? Where does the balance come from?

If you remember what we taught you empirical potentials, the way you set up an empirical potential model for rock sites is that you usually have, of course, the electrostatics, then you usually have an oxygen, oxygen potential-- because the oxygen is the biggest ion, they touch-- and then you have a cation, oxygen potential. Rarely do you have a cation, cation potential. So there's not even a direct potential between the two.

So if you actually froze these things on their lattice site, that V_{ij} is 0 because we just said that the electrostatics disappears. There is no direct cation, cation term, so V_{ij} is 0 in the case. So I sort of give the answer away and I said if you freeze these things on their lattice site, where does the interaction come from?

Well, if you freeze them, there's no interaction. So all the interaction comes from relaxation. OK? The interaction comes from the fact that calcium is bigger than magnesium and so that sets up an effective interaction. So it's kind of cool because with these interactions you have a Hamiltonian and a lattice model, which people always tend to associate with frozen positions, but the interaction comes essentially from the positions not being frozen, from the fact that the bigger atom pushes the smaller one away. OK?

And so this is maybe the key point to remember that once you parameterize your Hamiltonian with relaxed calculations, that Ising model represents the energy of those relaxed configuration. Not the ones where you're frozen at lattice sites. So when you do hydrogen on palladium, a lot of that interaction that you get between hydrogen on one spot on the surface and the other one has nothing to do with hydrogen, hydrogen interaction. It's actually all has to do with relaxation of the surface sites around the hydrogen.

You put a hydrogen here, the palladium around it relaxes. That's how it sees the other hydrogen, not because of a direct interaction. And with hydrogen that's not so dramatic. If you that with other absorbents like oxygen, it's very dramatic. They largely interact through a strain induced in the surface.

OK. So then you do Monte Carlo, and here's the phase diagram. It's a bit too many things on this slide, but essentially this is the line you should be looking at. This is the one with potentials, so these are the solubility limits at low temperature. So it's a two phase region in between, but you have some solubility of MgO and CaO and of CaO and MgO, and this is the quantum mechanical one.

So of course this becomes a miscibility gap because there's no liquid in the calculation so if you somehow could include the free energy of the liquid it would cut off that miscibility gap and form a [INAUDIBLE]. OK? But because you don't have a liquid, you're essentially looking at only solid state transitions. And again, you see the same thing.

The quantum mechanics gives you a higher solubility limit, so a lower temperature which you can get full solution. Because it's sometimes easier it can deal better with the lower symmetry state. OK. So this is the classic technique now to deal with any substitutional problem and I would say almost any crystalline problem because even if you work on systems where there are multiple lattice types in play-- I mean, if you think of simple metals you could think, well, maybe there's some that ordered on a BCC lattice and some that order on FCC lattice, you can cluster expand them separately and obtain free energies from Monte Carlo and then put them together.

That's why this is essentially the method by which any multicomponent crystalline problem is equilibrated, but I'll show you some of the older work just because historically it's kind of cool. This is the lithium aluminum system, which was I think only the second or third system that was done at [INAUDIBLE] by this method. This was done in the early '80s.

So the reason they did this in the early '80s because everybody thought we were going to make airplanes out of lithium aluminum because lithium is the lightest solid element in the periodic table. So any lithium you can put in an alloy makes it really considerably lighter, and the interesting piece occurs here, which you can see.

There's a delta phase here, which is metastable, sits about here. And what typically people do is they solution to heat up high temperature then [? quench ?] down so that you precipitate out delta phase to strengthen the alloy. It's difficult to do experimentally because this is a metastable phase, so it's a little hard to study it. But of course computationally, all metastability means is that there's something else that's more stable, so as soon as you remove that from the picture and the simulation you can stabilize it.

And so this was one of those early calculations where they did exactly what we showed here, build a cluster expansion, parameterize it, [INAUDIBLE] functional theory. In this case, did Monte Carlo do the statistical mechanics to find temperature and essentially get this calculated phase diagram? And this is the equilibrium phase diagram, but then you can also do the metastable. This is the delta phase precipitated by essentially removing the free energy curves of the stable phases and look at them. The liquid was kind of empirically added, I think, to this calculation.

This was my own PhD thesis, so I have to show it. Most of my PhD thesis were just what is 1, 2, 3, 4, lines, so I got off easily. But this is the phasing [INAUDIBLE] copper oxide, which probably most of you familiar is a fairly well known high temperature superconductor. And the critical issue, or one of the critical issues was, how the material behaves as a function of oxygen content.

Essentially that z up in the formula can go anywhere from 6 to 7 and at 7 it's a high temperature superconductor, at 6 and 1/2 it's a low temperature superconductor, and at 6 it's an anti ferromagnet. And so figuring out what the phase ion looked like was, I guess, worth the PhD thesis. And essentially it's, again, a lattice model problem. You have a bunch of sites where the oxygen can sit.

When you start taking out the oxygen it's now the site can be vacant or filled with oxygen. So it's a binary problem. So you cluster expanded the same thing. In those days, the quantum mechanics wasn't all that good and you did Monte Carlo and you get a phase diagram.

And what's neat about this O₂, to my knowledge, this phase-- I think there's just some history with it-- was probably the first structure I think that ever was predicted with quantum mechanics before it was seen experimentally. Now there have been a lot of cases of that, but this was actually in the late '80s, and it may have been the first case where the experimental evidence came afterwards. OK.

Here's a lot more complicated systems, so this is the lithium cobalt oxide I showed you in the beginning. There's cobalt layers, there's oxygen layers, but the key component is the lithium layer. As you take lithium out you have a vacancy lithium problem. So again, you expend in variables that only describe the occupation on this layer, OK?

So you don't really have to deal with the cobalt, with the oxygen, you just deal with that when you do the energy calculations and you get a bunch of effective interactions. And again, let me stress that these ECI are effective interactions. If you look at the first neighbor, it's only 30 milli electron volt, OK? To give you an idea, those first neighbor distances in the lithium plane are about 3 Angstrom.

If you evaluated the electrostatic source of a plus 1 charge and a plus 1 charge over 3 Angstrom, you'd get something that's close to an electron volt. OK? But you see, because you're taking effective interactions you essentially remove most of that. Well, first some of it is screened-- because you're not in a vacuum doing your electrostatic, some of it is screened by the oxygens that are close by-- but you're also canceling a lot of energetic effects and really seeing the only part that's relevant for the thermodynamics.

And then you can, again, do Monte Carlo just like we've done in the lab, get a phase diagram. This is the kind of phase diagram that you would not get out of straight Monte Carlo. The hydrogen and palladium problem we gave you, you can do that literally by looking at the output of the Monte Carlo, by looking at the heat capacity, looking at the energy, because the system equilibrates very well.

When it doesn't equilibrate you really have to do free energy integrations and look at where they intersect. This space also was predicted three years before it was seen experimentally, so this approach has a really very powerful predictive capability because you control almost all the approximations. So essentially, you have to live with the approximations in density functional theory. So whatever the accuracy of your electronic Hamiltonian, you can't do any better. OK?

But from there, you control all the approximations. So if you don't like the convergence of your cluster expansion, well, you add more terms to it. OK? So essentially you can make as accurate a representation of your quantum mechanics as you want. And that's why you can make Hamiltonians that are very accurate, but they live on a lattice model so they're extremely fast to search.

I mean, you've done the hydrogen on palladium. Hydrogen on palladium, that trivial job [INAUDIBLE], this 4, 5 million spin flips per second on a stupid, old IBM computer. So on a new machine, you can run 10, 15 million spin flips per second, so you can search configuration space extremely fast. OK? And so that's why it's so good at finding new structures.

OK, I was going to show you a surface problem. So here's platinum 111. Reason I took platinum, it's sort of an important material for all kinds of catalysts and so. And every time you have these FCC surfaces, the possible absorption sites essentially form a triangular lattice. I should go to a different color here. So you have these really old form of triangular lattice.

There we go. Triangular lattice. So if you want to study the absorption of something like hydrogen or oxygen, again, you build a lattice model. Then you have to do density functional theory calculations to parameterize it and I think you've done that now in the lab. So because this is surface you would set up a slab geometry. So you essentially would take a few layers of platinum and then put stuff on the surface to absorb it, and you arrange a vacuum above it and do the DFT calculations.

OK. So here's that result. Again, this shows both LD [?] and GGA [?] but again, we plot them-- this is oxygen absorbed on platinum 111. We reference the full coverage and the zero coverage to 0. You may be a little surprised that you can set two points to 0 on the energy scale, but one, of course, is obvious. So you can always reference one point to 0 on energy scale.

The reason you can set the second one to 0, what you're essentially doing is setting it 0 for the chemical potential because the chemical potential is, in some sense, the slope of that curve. It's the energy changed just by adding oxygen in this case. So you're essentially referencing the chemical potential to 0 by setting two points to 0. So you do that, same thing. You fit interactions.

I've actually made some copies of real notes so you see how this is really done. So people sketching out on triangle lattice, different interactions. You find their value, again, they're very small. The black points is for oxygen or platinum, the other one is on ruthenium. You know it's about 75 for a near neighbor interaction and then it decays very quickly very fast.

And you know that they actually have to be in that range because if you're going to see phase transitions in the range of, say, 0 to 1,000 Kelvin, your energies have to live on that energy scale and 1,000 Kelvin is what? About 100,000-- well, 1,000 Celsius is about 100 milli electron volt so your interactions have to be in that range for stuff to happen. So I thought I'd show you, rather than a nice picture of a phasing-- hold on.

So how would you look for phase transitions? You do the same stuff you do in the lab, you track heat capacities, you track energies. So this is for the same run, different oxygen chemical potentials. So with your experience, what's happening here at about 940k? This clearly looks like a phase transition, first or second order.

What's your guess? Second? It's hard to say, but it looks like second order. I mean, it's maybe not refined enough to look, but there's clearly a peak in the heat capacity. Let's say we take, I don't know, the green curve, which is at the particular oxygen there's a peak in the heat capacity. It's not very well defined, you shouldn't let yourself mislead.

If you didn't draw the lines here but just the points, that peak is essentially defined by one point. I mean, you should always be careful. You may have very spurious behavior going on in one similar Monte Carlo solution. The other peaks have slightly more points, but normally if you really wanted to know which is first or second order you would refine it a lot more near the phase transition.

Again, the energy seems to be continuous, but don't let the lines fool you. You've taken, say, this point and that point and connected them, so of course they're continuous, but maybe what this thing does is it continues and then shoots down to the other vertically, so it could still be first order. So you would normally refine this a lot more.

And the other thing, if you really want to know that something's second order, you look at the scaling of the heat capacity with system size and that's the really accurate way to know that something's second order. You're normalized heat capacity in a second order transition, so the heat capacity per site scales like the square root of the number of lattice sites, OK? You'll find that already in your simple example, your simple homework problem of hydrogen on palladium, you'll find it remarkably hard to identify in certain regions whether the phase transition is first or second order.

That's why you put them in, they are both first and second order. So [INAUDIBLE] beautiful phasing, I thought I'd show you guys how it's really done. When we do Monte Carlo, this is really like a puzzle, you know? In papers we always show these nicely drawn things, but this is how you get to it. You do an enormous amount of scans in μ . So you scan this way, then you scan in temperature, and you slowly start piecing together which phases exist where, where they disappear, how they disappear, first or second order reactions.

So this is one that is near completion with a few question marks whether these are like first or second order. And then later you draw nice lines, but this is one of the problems with Monte Carlo because it's a stochastic method. It really is a lot a lot of puzzling. And what you'll find is that the last thing that's useful, the least useful thing is actually looking at the spins. That's people's first intuition.

I'm going to actually look at the configuration in my system. You know, any time systems are partially disordered, just looking at them is the best way to fool you. If you look at something partially disordered, half of the class will say it's ordered and half of the class will say it's disordered, and it's really the thermodynamic behavior that tells you in which phase it still is.

If it's in an ordered state, so in a compound state, that means that you should be able to push it to a disordered state. OK? That means you should see somewhere a phase transition as you raise the temperature or drive the chemical potential to another phase. If it's already disordered, if you heat it further up it won't go to a phase transition anymore. OK.

Just to show you that you can push this a lot harder and make it a lot more complicated. I'll show you something that gets a little more abstract, and it's from a problem that I worked on quite recently. It's absorption induced segregation, and here's the problem. So the problem before was oxygen absorbing on platinum, but if you have actually a platinum metal alloy like a platinum metal catalyst, in this case, platinum ruthenium, then you have two configurational problems.

So the first question is, what's on the surface? Is it oxygen or vacancy? Is the site oxygen or vacancy? But there's also segregation thermodynamics now. On the surface, you now may not have pure platinum, you may have either platinum or ruthenium and the two interact. What segregates to the surface is influenced by what absorbs.

You see, you may think of that as a quaternary problem because essentially, what's in play now? Platinum is in play, ruthenium is in play, oxygen's in play, and vacancies are in play-- vacancies being vacant absorption sites. But it's actually a 2 plus 2 problem. It's really a binary configuration problem in the surface layer and then a binary configuration problem on top of the surface layer, and so you can still perfectly solve this with a binary cluster expansion.

It's just that your spin means something different whether you're in the surface layer or whether you're in the absorbent layer. So you could call them different, but you don't have to. You could say, well, there are variables that is, say, this describes what's in the absorption layer and I could just call a new set of variables delta, this describes within the surface layer, and then I could literally expand in the products of these. But to be honest, you could just all call them sigmas. OK.

They just mean something different whether they're defined on this site or on that site, OK? And so you can do the same technique and now essentially expand a quaternary problem. And you do the same whole thing, you calculate a lot of configurations. Because your configuration space is essentially bigger, you'll need more information to parameterize it. And then you do Monte Carlo and equilibrate it and here's an example.

This is, at a given temperature, driving up the oxygen chemical potential so we expect more oxygen to absorb, and that's what you see. But you can actually track the amount of ruthenium segregating to the surface and as you bring more oxygen to the surface, you actually start increasing the amount of ruthenium in the surface as well. And so you can really solve the thermodynamics of these problems by [INAUDIBLE] initial methods by going through this cluster expansion, but you can actually look at the sites.

Orange is the platinum surface, red is the oxygen absorbed, and this picture here is the same with the oxygen stripped off so you see underlying ruthenium. So as you increase the oxidation, you essentially get more ruthenium to the surface and more oxidation. This is sort of more and more, higher and higher oxygen pressure. OK.

Let me show you one more quickly before I continue the theory and the reason I want to show you this is because it's sort of a little different. This is a structural problem from hydrogen modified aluminum fracture. OK. The problem there is that impurity embrittlement in metals is a big issue.

There's a lot of metals whose ductile behavior becomes very poor when there's impurity present, and hydrogen embrittlement in particular is a bad one. And it's not that hydrogen's so much worse than other impurities but first of all, hydrogen tends to be a common impurity because it's available everywhere. Any time you have corrosion reactions going on, you often produce hydrogen.

Any time you're working in acid environments, you produce hydrogen, and it tends to be a fast diffuser. So there's a lot of other impurities, but a lot of other impurities sort of never make it to the problem areas in the material but hydrogen tends to be a fast diffuser. So the way people have dealt with this in quantum mechanics is literally by looking at how planes decohere here with and without hydrogen.

And of course, if you hydrogen in between it, it decoheres more easily. There's less work of decohesion. But if you want to look at the structural problem of what happens when you do this with hydrogen, if you want to equilibrate the hydrogen it becomes a lot harder.

So let me sort of first tell you how would you do this without hydrogen. You could make a slab, you would define a plane where you open up the slab and you remove the two slabs farther and farther apart so it's like you're opening a crack. You keep track of the energy and the derivative of that is, of course, the force required to separate the slab. But think of the problem now if you want to do that with an impurity.

You have to now equilibrate the impurity at every separation of the slab. OK? So you pull the crack open a little more, you now have to figure out what happens to that impurity there because as that impurity equilibrates, not only its structure equilibrates by the way, but also its composition because segregation to those surfaces is an open system problem. The concentration there is not fixed because it comes from the bulk or from a gas phase.

So it's actually a grand canonical problem, it's a constant chemical potential problem rather than a constant composition problem. So you see that the hard problem you have. Not only do you have to arrange the impurities in the most stable way because that's going to determine the energy and therefore the force it requires to open up those planes, but you also need to equilibrate the amount of impurity in that crack.

And so this is actually why it's so hard to include chemistry in mechanical calculations, because you're really working on two different time scales that are very hard to combine. And so the solution here was essentially to calculate the relevant free energy of hydrogen for any state of crack opening, for any state of separation of the planes you tried to get the equilibrium hydrogen content by a cluster expansion and then from that envelope of free energies you get the force.

So it's again, the same thing. You define a lattice model. This was hydrogen on one side of the crack opening, hydrogen sites on another side of the crack opening. So now you have interactions within one plane, but also between the two. You do a bunch of first principles computations, so it's, again, the same thing to fit those interactions. Then you're going to do Monte Carlo. You're going to, in this case, get your free energies in the derivatives of the free energy with respect to the separation is going to be the force.

So in this case, this is the force as a function of slab separation for all the different hydrogen coverages. So I plotted the force. I really should have plotted the energy because the energy is the quantity that's expended. And then you get interactions. In this case, I'm only showing two.

So there's an in-plane hydrogen interaction, that's the red one here, and there's an inter-plane. The inter-plane is between the two openings of the crack and of course, that one tends to go to 0 as you separate more and more. OK? So now the cool thing is at every state of the crack opening, you can actually get the free energy of the system fully equilibrated. And so if you get the free energy, you can actually find the force, and this is the fully equilibrated thermodynamic trajectory, that red line.

And what's really surprising is that because you fully equilibrate the system, it suddenly looks very different from what you thought it would be and what you would have gotten in an unequilibrated calculation. These black lines, this is what you got in an unequilibrated calculation. It sort of makes sense, you know? As you pull, the force goes up. No surprise, I mean. And then you reach a critical point, a point of maximum force. At that point, you're unstable because essentially the force, as you pull further, goes down so you're unstable.

On the downslope you're essentially unstable. But look at the red line. So this as a chemically equilibrated system behaves very different. Essentially you pull a little bit, you open the crack a little bit, and then the crack opens with no increase in force. OK? And then you pull the saturated crack apart. And essentially what this is, because you've done your thermodynamics right, you actually have a first order transition in length space.

OK, this is a length discontinuity. It's just like a volume discontinuity in a regular first order transition, OK? Here, essentially the crack is separating at constant force. OK? It's like when you do a first order transition. Upon heating, you're putting an entropy at constant temperature for example, or if you do a pressure induced phase transition, you get a volume change at constant pressure. OK?

So it's the same thing. But you see that by doing the chemistry right and the equilibration, you get very different results. OK. So the last thing I want to do is come back to the idea that we started with and I wrote down the energy as a function of configuration and then showed you how you cluster expanded.

What I want to do is now come back again to what the definition should be of that energy, and we did that at the back end of the Monte Carlo lecture but because it's kind of important I want to come back to it because it's not totally clear what that energy or that Hamiltonian in the lattice model should be. So let me take a step back-- and I know you've seen this, but I'm going to go through it.

If you want to get the free energy of a system you normally want to include all excitations because essentially you have to calculate the partition function, you have to trace over all excitations, and that will give you the full partition function and the [INAUDIBLE] will give you the free energy. So if you think of the most basic system, we'll always have vibrational excitations, OK?

If we have anything with a binary problem like A's or B's or something and vacancies, we will have configurational excitations, and then unless you're in an insulator you'll have electronic excitations. The idea we talked about last time is that the way you can define a Hamiltonian that lives in this space, which is the one with the slower time scale, is by integrating over the faster timescale.

And I started to show you last time I think when we finished the Monte Carlo lecture. And let me sort of go through that formal description again so we see exactly what our latest model Hamiltonian is. So I'm going to do it here only with vibrations of the fast time scale and substitutions as the slow timescale, and after you'll see how it would go with electronic excitations as well.

So I replace the coordinate r_i , so this is a set of coordinates that describe the atom, into a lattice site position and a deviation from that lattice site. OK? The lattice site position, that is fully defined by the set of configurational variable, σ , and the deviation from the ideal lattice site position are called configuration ν . So the combination of σ and ν is really the same as the set of vectors r_i , OK?

So what I'm going to do is-- you already see where I'm going probably with this. I'm going to trace over the ν s to get a Hamiltonian in the σ s. OK. So the full partition function is the sum of the exponential of the energy, which depends both on σ and ν , and I sum over the σ states and the ν states. OK? What I'm going to do is I'm going to do the inner sum so that I'm only left with the sum over σ s of $\sum f$, which is the Hamiltonian in the configurational space.

And if you look to make these to the same, what that f should be? Well, f is the free energy that I get by tracing over the ν s. OK. I mean, there's no physics here. This is pure and simple substitution of variables. So I call, essentially, this sum here is the exponential of minus βf , which means that f is the log of that exponential.

So look at what I formally end up with. I have a partition function that only sums over substitutional states. So this is a problem that lives in the same space as a lattice model because I'm only coming over discrete configurational states. So what I'm really saying is that if I do this partition function correctly, I have the full partition function and therefore, the free energy.

But to do it correctly, my Hamiltonian in my lattice model state needs to be f , not e . And what is that f ? Well, f is essentially a free energy. It's the log of a partition function. And what partition function is it? It's the partition function by summing over all the faster states. So the problem is, essentially, simple. The Hamiltonian you should use in your slow time scale space is the one you get by integrating or by calculating the free energy of your fast time scales.

So in this case, what should you do? In everything I showed you earlier this morning, I used the energy as the Hamiltonian. OK? I use the direct to DFT energy in my Hamiltonian. What I really should have used is the vibrational free energy to be correct. If I use the vibrational free energy in that Hamiltonian, I would have gotten the full partition function and therefore, I would have gotten the total free energy of the system.

So what did I do instead? I didn't calculate this. I approximated this by something. I approximated this by the state in this ensemble with the lowest energy. I mean, if you think of this as a free energy, this is going to be a free energy plus some entropy term, which you could think of as a ground state energy plus contributions from excited states. OK?

I've basically said, I'm taking this just as the ground state energy in this [? ensemble ?] because I'm taking the lowest energy configuration that I can get by varying over this. OK? And what is varying over this? What are the n 's? The n 's are the displacement from the lattice sites. So essentially I'm relaxing my system, letting it displace from the lattice site, I get the lowest energy state in the ensemble, and that's what I expend.

So when I'm done, what do I have? When I'm done and I do Monte Carlo with that Hamiltonian, I have essentially the right energetics, I have configuration entropy, but I don't have vibrational entropy. To have vibrational entropy, I would have to do this sum, not just take the lowest energy term. OK. Come on.

OK, let me sort of skip that. OK. People actually have done examples where they do the full integration over vibrational states rather than just approximating that ensemble with its lowest energy term. You see that it's a lot more work. Remember that to parameterize the cluster expansion, I need the Hamiltonian evaluated for a bunch of configurations. Even in the simplest case that's easily 40, 50 configurations.

So if you want to include the vibrations, you actually have to do, for 50 configurations, the full phonon DOS or do molecular dynamics on them and integrate the phonon and get the vibrational free energy through that. So it's like two orders of magnitude more work. So rather than just the energy, you want to get the whole vibrational spectrum, and that's why people usually don't do it. But the nice thing is that in this approach, you can do it if you really want to.

And I'll show you one example where people did it and looked at the difference. It was the very simple case of the cadmium magnesium phase diagram. It was done by Mark Asta, who's now at Northwestern. So this is the experimental phase diagram actually. So this is the calculation with no vibrational entropy and with vibrational entropy. So the first one is the one that I sort of showed you as the standard one. You just take the lowest energy for a lattice model configuration and you do Monte Carlo then with that.

And this is if you literally do the whole phonon DOS for every configuration. And what you see is that there is some change in topology. I mean, look at this point has come down considerably and even has changed the topology somewhat, and usually that's what you see. You rarely see new phases new ground states come in when you include vibrations, but you will change the topology of phase diagrams, for example.

There are a few known cases. Usually people write a physical review letter whenever they find one where phases come in purely stabilized by their vibrational entropy. It's a bit of a puzzle, actually, why it doesn't happen more because the vibrational entropy is large. It's essentially $3r$. Remember the law of the [INAUDIBLE]? The entropy is $3rt$. The [INAUDIBLE] is $3r$, so the entropy is $3rt$.

The configurational entropy at its maximum is only $r \log 2$. At 50-50, in a binary alloy, you have the maximum entropy and it's actually the logarithm of $1/2$, so it's minus the logarithm of 2. So you either have-- that's about, what, 0.67? So the configurational entropy only about $0.67r$, the vibration one is $3r$, so it's a lot bigger.

But of course, what matters is the difference in that quality between different structures and different states. And even though the configurational entropy is the smaller one, it tends to be the one that's more different between different states than the vibration one, and that's why we sort of all hope that it's OK to often leave the vibrational entropy out because often most of it tends to cancel between different states. OK.

So sometime in 1994, somebody compiled all the systems done with these methods. This was 1994. We stopped compiling after '94 because that was just way too many studies going on, but this kind of approach has now been used for the study hundreds and hundreds of [INAUDIBLE] structure and thermodynamics problems. People do higher component now, ternary's it's still a lot more work. This is a simple ternary.

And then I gave you a series of references. I wanted to show you some of the older work, that's actually the last references here. Often the older work, even though it's not up to date, but it's also not confused by a lot of complications so it's often very easy to read, very elegant to read. But the key review paper is essentially this one, which is like 150 pages or something like that which contains the ideas of the cluster expansion and a lot more.

And then my own work is just a short review on how that works, how you apply the [INAUDIBLE] sites, and the one by Alexander Zunger is essentially the same thing but how you apply that in semiconductors. And the fourth one is the most recent one, which my own grad student wrote, Axel Van de Walle, because he's been kind of instrumental. He's gone to making this into automated codes and so on, has a whole code suite and the paper is about that, how you train Monte Carlo and so to automatically look for phase transitions and things like that.

OK. I think that's it. So let's see, what day's today? Today's Thursday, so Tuesday we'll meet again here. I'll start talking about accelerated molecular dynamics and find another topic to talk about as well, and then Thursday is Chris Wolverton. So I'll see you, have a good weekend.