PROFESSOR: --called case studies. That is, we really look at how some of the tools that we have learned during this class-- some of the energy tools, like density functional theory, some of the statistical tools, like thermodynamic integration-- come together and are actually used to predict the properties of materials and structures of real interest. This is-- I called it lecture 20, but I think I have been slightly inconsistent with the numbering of my lectures. We count also the labs. Anyhow, everything is on the web.

So what will you learn today is actually how to become an expert in two of these jobs for the 21st century. One is the virtual alchemist, and the one that is fancier, I think, is the nanotechnologist. There was probably an ad for GE a few months ago on TV in which there was a nanotechnologist marrying a supermodel. It was sort of fairly intriguing for all of us doing the initial calculations.

But more seriously, what we are really going to focus on is how can we break the n-cube barrier of electronic structure calculation. You have learned that, say, when we study a system using density functional theory, the computational cost scales as the cube of the number of atoms. And that basically scales very rapidly the size of the system that you can do. I mean, even on a supercomputer nowadays, we can't really do more than 500 atoms, say. And on a regular cluster, we are more likely close to 100 atoms.

And the reason of the cubic scaling in DFT is ultimately due to the orthonormality constraints. Every single particle orbital needs to be orthogonal to every other single particle orbital. So if you double the size of your system, you double the number of electrons.

That is, you have twice as many orbitals that need to be orthogonal with twice as many other orbitals. So you have four times more matrix elements that you need to calculate. But also, those matrix elements are calculated in a unit cell that is twice as large. So there is another factor of two, and that gives us a factor of eight.

And of course, there is a lot of effort in the community nowadays to go towards linear scaling algorithm, algorithm that just double in computational cost when you double the size of your unit cell. And they are all based on the sort of physical principle that, in many ways, nature is local. So your electron needs only to know what's happening nearby and doesn't need to be really orthogonal to electrons that are very, very far away. And you see an example of sort of where this linear scaling direction are heading.

Other methods scale even worse than cubic. We have sort of discussed that sort of standard Hartree-Fock scales as the full power, and highly correlated quantum chemistry methods scale as the fifth, sixth, seventh power. An interesting exception, I think, among the methods that try to go beyond DFT in accuracy is what people call quantum Monte Carlo approaches. They tend to be approaches that are very accurate but based on a stochastic sampling of the wave function, so really, a stochastic solution for the wave function in the Schrodinger equation.

And those in principle scale as the cube. And there have been sort of recent progress in making them truly linear scaling. So I think they are a very promising avenue for the future.

But in this class, we'll sort of confine ourselves to DFT. And in particular, we'll look at two examples in which we use electronic structure DFT calculation to actually obtain the relevant parameters for a Hamiltonian that basically models the important degrees of freedom for our system. So we'll actually look at semiconductor alloys, where you'll see, very naturally, one can sort of solve the problem of substitutional disorder on that alloy using an Ising-like Hamiltonian, so something that you have seen in the previous classes.

And then we'll look at another example to functionalize nanotubes, that is carbon nanotubes, to which organic ligands have been covalently attached. And you'll see how we sort of develop a linear scaling tight binding-like procedure. We'll explain this. That allows us actually to study, with full DFT precision, systems that have tens of thousands of atoms in the units cell.

So let's start with the first example, that of semiconductor alloys. Semiconductor alloys have been of central interest since, I would say, the late 80s, when people started to realize that there wasn't only silicon. And in particular, it was very important to tune out the physical properties of your semiconductor, and alloying is a very simple way of doing that. If you, say, mix together something like gallium arsenide and aluminum arsenide, depending on the concentration, you have a substitutional alloy with a bandgap that is variable. So it can be used very efficiently, say, to tune the emission frequency of your laser light, if this semiconductor is your active layer.

And this is something which, in our department, in the material science course, there is a lot of ongoing effort. In particular, Professor Fitzgerald works a lot on some of these alloys to actually improve mobility of the carriers.

And I had here some kind of beautiful picture of actually an array of vertical KVT lasers in which sort of a number of vertical pillars were grown on a surface with a composition gradient for that active layer. So what you achieve on a single breadboard is a very high density of emitting lasers. And each one emits at a slightly different wavelength, because the bandgap is tuned by the compositional gradient in that array.

And another set of reasons that sort of intrigued and prompted the study of the semiconductor alloys was actually what is called the ordering, multilayer ordering, superlattice ordering in those alloys. That is, when doing X-ray diffraction of a standard semiconductor, you would get a set of broad pixel. And what people note is that under certain growth circumstances, they would find naturally in the diffraction figures some intermediate peaks.

So in this case, some peaks are closer to the origin, or in this case, sort of the green dots. So yeah, midway between the two red dots. What it means is that there is actually some kind of ordering that has periodicity that is in that particular crystallography direction, has a wavelength that is twice as large or a wave vector that is one half. And so it sort of hints at some kind of ordering. So there were a number of reasons why these were interesting. And I'll show you sort of one of the ways one could actually tackle the problem of studying these systems.

I had another figure here. They just didn't come-- they didn't come out-- in which, again, I was sort of showing you basically the whole set of semiconductors and plotting that as a function of lattice parameter and bandgap. And basically showing that, if you look at all, say, the III-V semiconductor, they sit in a certain region of this diagram. But in reality, there isn't anything that is really useful that-- say, as the bandgap of silicon. But it meets at the wavelength for which conventional fiber optics cables are transparent. And so a lot of effort in trying to engineer alloys or new materials that go towards that soft spot in which you can grow things that seem more thick on silicon. So you can use all your silicon technology, but they have the right direct bandgap that is appropriate to optical fibers, 1.5 microns.

OK, so from the conceptual or computational point of view, our problem is really trying to figure out a way to calculate efficiently the energetics and the phase ability of a semiconductor alloy. And I sort of shown you here an example. Remember, semiconductors have the diamond or the zinc blende structure. So they are made by two, say, interpenetrating lattices. And what I'm showing here is just one of the two FCC sub-lattices of a typical zinc blende structure.

Say, suppose that you are studying gallium aluminum arsenide. Arsenic is the group V semiconductor. Gallium and aluminum are group III. So you would have a sub-lattice where gallium and aluminum sit, and this is what you are looking at. It's an FCC sub-lattice-- and where there can be substitutional disorder.

On each side, there is a certain probability of finding either a gallium or an aluminum that is basically proportional to the concentration. So if you have a 50/50 alloy, you will have 50% probability of finding gallium in a certain side. And then there is a second sub-lattice that is almost invisible that would be given by the anion or the arsenic side.

Of course, one could sort of study this system with a sort of brute force approach. We could really look at large supercell, and perform many calculations, and find out which arrangement of atoms give us lower energy configuration, figure out what, say, would be the bandgap or what would be the lattice parameter of certain configuration, and try to figure out what are the macroscopic properties. But statistics at this stage will always be limited. If you want to study a supercell that is large enough, you probably need, again, a cell in which the correlation between faraway atoms is lost, and you might easily need something like 100 atoms.

But then in particular, you need a lot of statistics that are true to the 100 ways of arranging atoms on that lattice. And you can only calculate a few of these. And so you become very limited in analyzing the thermodynamics of your problem.

At the end, what you really care about is understanding what is the energy of a configuration. That is, what is the energy as a function of-- you have seen this before-- configuration of variable sigma. We use the nomenclature in which sigma is equal to 1, if, say, we have a gallium atom here.

This example specifically is actually for gallium indium phosphide, so we'll have, again, group III indium. And sigma equal to minus 1. That means that there is indium on one side. This will make actually the connection with the Ising model very, very apparent.

And so what you want to do is calculate the energetics of a lot of configuration and then extract the thermodynamic properties. And you have seen already in Professor [? Seder's ?] lectures, and you'll see again partially in this lecture and the next lecture how you can use actually Monte Carlo approaches to extract basically what are the average ensemble properties. But our first goal for us is actually finding out an inexpensive way to calculate the energy of an arbitrary configuration of atoms, without having to do the explicit first principle calculation.

And this is where the fundamental concept that you'll see in this lecture comes about. That is, how we actually devise from the initial calculation, how we extract from the initial calculation, the parameters for the model Hamiltonian that describes the energetic of my system for any arbitrary configuration. And as usual, you need to put ingenuity in solving a difficult problem.

And this is one of the possible solution. You'll actually see some of different alternative approaches in one of the next lecture. But what we are doing here is, we are actually sort of exploiting the fact that, say, if you are looking at a substitutional alloy, you have substitutions on a sub-lattice between atoms that are really chemically very similar.

If you are studying, say, gallium indium phosphide, gallium and indium are going to be chemically very similar. They are both group III. They are just sort of in a slightly different in shape. And of course, their orbitals will be sort of slightly different, and they will correspond to a different period in the periodic table.

But their chemical properties are still fairly similar. And so for a moment, we could actually think-- and I'll show you how it works mathematically-- that a real configuration of a crystal-- and remember, from the point of view of density function theory, having a certain arrangement of red and blue atoms means that there are very specific pseudo potentials sitting on each one of the sides. But one of these arrangements could actually be sort of considered as a perturbation on what we call a virtual crystal, a crystal that has full translational symmetry. So the sort of idea that I'm trying to give you-- and we'll see it mathematically in a way-- is that we could think at the real configuration of a system as a perturbation in which we take a pseudo atom that doesn't exist but is somehow very similar to both gallium and indium.

And then we turn that with a small change in the pseudo potential in either gallium or indium. And if we can do that, well, then we can study this problem with density functional theory, because this is actually a system that has two atoms per unit cell. And then we can try and use perturbation theory to figure out what would be the energy of our configuration over there.

And this is actually fairly trivial to sort of see it in mathematical form, when you actually think again at the fact that in density functional theory, the potential acting on the electron is really given by this array of pseudo potential sitting on lattice sides. And so if we have a real configuration-- remember, the red and blue crystal on the left-- well, to that, it will correspond an external potential that is an array of red and blue pseudo potential, that can actually be rewritten in terms of the semi-sum and the semi-difference of pseudo potential. So this would be the exact expression.

We have the true pseudo potential on each sort of FCC sub-lattice site. But we write, say, gallium pseudo potential as being given by the sum of the average between the gallium and the indium pseudo potential, say, plus one half the semi-difference. You see, if I'm sitting on a site, and I'm telling you that the pseudo potential acting on that side is going to be the sum between the average between gallium and indium plus one half the semi-difference, you see that the indium cancels out, and the gallium sums up. And what you have is the gallium pseudo potential there.

And so a sigma variable equal to plus 1 at the site R really gives you a gallium pseudo potential in that site. And a sigma variable equal to minus 1 gives you an indium pseudo potential in that site. But in this form, and in the limit in which the two pseudo potential are fairly similar, you can see that this term is actually small. The more similar gallium and indium are, the smaller this system is.

But this term here overall is the only one where the configurational variables are present. There are no configurational variable here. And so this represents, really, the green average virtual crystal to which at every site, we add or subtract depending on the sign of sigma R, the semi-difference of the two pseudo potential. And this term really turns a green atom in a red or blue atom.

And for the specific case of semiconductor alloys, this term is smaller, because disorder takes place between atoms that are very small. And so what we are going to do is use density function theory to study this crystal then use perturbation theory to understand what is going on when we do this perturbation. And so this is also sort of a chance for me to at least hint at how we actually perform perturbation theory in density functional theory itself.

And that's a very useful tool. Say, it's used, and you'll see it, and you have seen it before, to calculate things like phonon dispersions. That is, what is the force or what is the frequency of perturbation in the atomic position with a certain wavelength that is really the normal mode of the crystal. Or it can be used in a variety of problems, say, to study what is the response of a system to an electric field and then calculate Raman or infrared spectra.

It can be used to calculate the response of a system to a magnetic field and then calculate the NMR spectra, or it can even be used to calculate higher order important physical quantities. If we go to sort of higher order response, we can start to calculate what is the third order Hamiltonian in the phonons, so what are the-- in the phonons, what are the lifetime of phonons. Or we can calculate how an electronic state couples with a phonon state. And so we can actually get out the parameters, say, for superconductivity. That is, when electrons travel through a crystal paired because of the interaction with the phonon degrees of freedom. Or we could calculate, say, resonant Raman scattering that, again, is sort of a strong electron-phonon component.

But in general, the problem is just this. Perturbation theory looks at what happens to a certain system that has an external potential v0 when you add a perturbation. And the perturbation is sort of parameterized by a certain strength lambda.

So for small lambda, the perturbation is weak. So we are sort of thinking at the limit of what happens when the perturbation is weak. So you add a lambda delta v term to your external potential. And in this specific case, that delta v term is going to be that semi-difference between the pseudo potential. So it's the term that really transform a virtual atom in a real gallium and in a real indium.

What happens when you apply this small perturbation is that, of course, your new self-consistent charge density will change. So if, for a given external potential v0, you add a ground state charge density n0 for a perturbation in which you add the term lambda delta v, your new ground state charge density is going to be a certain charge density that will define us n lambda.

But we can think of this n lambda as being sort of expanded in a power series of power lambdas. So the new n lambda is going to be given by the original ground state charge density plus a term that we call the linear response that is a first order in lambda, plus a term that is a second-order response, plus a term that is third order and fourth order and so on and so forth. And the smaller the lambda is, the less relevant will be the higher order terms. And what is really important is just this first order response in the charge density.

And this actually sort of gives us a lot of power. That is, just knowing the first order response in the charge density allows us to calculate the energy of a system in the presence of this perturbation up to the second order. And this basically can be seen by using Hellmann-Feynman theorem here.

This is-- you have seen it sort of when we discussed forces, and it's really a sort of direct consequence of the ground state properties of the orbitals in the Schrodinger equation. But basically, what does Hellmann-Feynman theorem tells us? Well, it tells us that if we need to look at the change in energy with respect to a certain perturbation, the only thing that we need to calculate is actually the expectation value of the orbitals with respect to the perturbation. So this is really Hellmann-Feynman theorem.

You might remember this. We have discussed it. We don't need to take into account the derivative of, say, lambda-- sorry, I'm using too many lambdas here. But let's call it lambda again here.

We don't need to take into account the changes in the wave functions themselves. We just calculate the expectation value. So for any given arbitrary lambda, the derivative of the energy, the change in the energy with respect to lambda, is just given by this simple integral.

And so you see that if we then say that we actually expand the energy up to the first-order term, what we have is that the energy is equal to e0 plus lambda d e over d lambda. And then there would be a one half lambda square term. But then in itself, we can expand-- it's a little bit father, so I'm sort of skipping some of the passages-- we can expand this n lambda in powers of lambda.

And so what you see is that up to the second order in lambda, the only terms that are relevant in the expansion of your energy-- and this is what is important-- are, for the first order term, just the integral of the perturbation with respect to the ground state's charge density, and for the second-order term, the integral of the perturbation with the first order change in the charge density. So perturbation induces a change. And if we are able to calculate the linear response to that, this n1, we actually get correction to the energy up to the second-order term.

And so there is a whole branch of density functional theory that's called density functional perturbation theory that really deals on how to calculate this linear order perturbation. And I really won't go into the detail. The sort of other difficulty with respect to a standard set of quantum mechanical perturbation theory is that in density functional theory, the Hamiltonian is in itself self-consistent. And so you have an added layer of complexity, because it depends from the charge density itself.

And so when you have got a perturbation in your external potential, the perturbation in your Kohn-Sham potential contains also term in the Hartree and in the exchange correlation term. So we won't really go into these details, but there is, again, an iterative self-consistent procedure to obtain the self-consistent linear order response to an external perturbation. And all of this is actually implemented in the PWSCF code. So some of you might at certain point need to use that, again, to calculate quantities like the phonon dispersions. And there are actually examples in the distribution that tell you how to do that.

And so if we go back to our original problem, all of this was actually to convince you that if we are able to calculate this linear response, the first order change in the charge density induced by a certain perturbation, we actually have the energy of a system up to the second order in the perturbation, where the first order term, again, is just given by the integral of the ground state charge density in the perturbation. And the second order term is given by the integral of the linear response times the perturbation. But you see, what is now the fundamental step is that these quantities don't depend on the configurational variables. These are quantities that are calculated on the ground state-- that is, the green crystal-- something that doesn't depend on configuration.

So this can be calculated once and for all just on the system that contains two atoms per unit cell. And once you have calculated this, this will give you the energy up to the second order in the configurational variables. So doing this calculation will take some time, but it's really on a two-atom unit cell. But once you have this, what you have is really a analytical and extremely inexpensive expression to calculate the energy up to the second order. So as long as your disorder is weak, you will get a very good approximation of the energy just by doing this set of sums that a computer will take probably a nanosecond or something like this.

And again, this really looks like an Ising model Hamiltonian. If you remember what you had in the Ising model-- is that you had a sort of coupling between nearby spins that was parameterized by a certain exchange parameter here. The only difference is that, being these are semiconductor alloys, interactions tend to be long range.

And so it's not only atoms or sigma variables sitting nearby that sort of interact, but the interaction sort of goes on farther and farther. And so it's a sort of long range Ising Hamiltonian. But once you have this, you can do all the thermodynamics that you want to your problem and very inexpensively.

I'll sort of go very quickly through this. But again, what we have seen up to now is really this sort of fundamental concept that the electronic structure calculation or our energy calculation allows us to extract parameters-- in this case, the J interaction constants-- that gives us the energy of our system. And as usual, when you really need to find out what are the parameters to extract, you need to think at what are your relevant degrees of freedom. So you need to think at the physics of your problem.

And if you have a semiconductor alloy, the first degree of freedom that comes to mind correctly is the configurational degree of freedom. That is, we can have on a certain site one atom, gallium, or another atom, indium. But actually, it's not the only degrees of freedom that is relevant, because depending on which atom you put on a certain side, what you have is that, depending on also what is the surrounding environment-- if it's a gallium between gallium atoms or if it's a gallium between indium atoms-- that atom will see a broken symmetry. And so it will want to relax from the ideal sort of periodic Bravais lattice position and will likely like to move around, either to make space for larger atoms that are nearby or sort of to make space for itself if there are smaller atoms.

And so this is sort of where intuition becomes important. And so you need to decide what are the important degrees of freedom. And in this case, what we are saying is that there are not only configurational variables that are important degrees of freedom, but there are also displacements of the atoms in response to a certain configuration away from equilibrium.

And so once you have identified that there is this second set of degrees of freedom, well, what you are really saying is that we really need to consider, in the expansion of the energy over all my degrees of freedom, not only the configurational variables but also the displacement of an atom, say, from its equilibrium position. And so we need to expand their energy up to the second order not only in configuration but also in displacements. And because we are expanding up to the second order in configuration and in displacements, there is going to be also a mixed term that is half configuration and half displacement.

If you want this term, think of F as being the second derivative of the energy with respect to u and with respect towards sigma. So this term here is telling us what is the force that is induced on the atom R when, on the side R prime, there is a certain atom given by the configurational variable sigma. So we really need to expand in all our degrees of freedom.

But then we can do something more. And we can even sort of apply an equilibrium condition. Because at the end, what we want to do for the system is not really the dynamics, but what we really want to do is the thermodynamics of the different configurations.

So what we are saying is that, well, suppose that we have chosen a certain configuration of atoms. Now, what we really need to take into account is not only how the energy changes but also how the atom moves around to accommodate the configuration as best as possible. But that really means that the atoms will move in order to minimize the energy, or if you want, in order to have really zero force at a certain point.

If you start from the periodic lattice and you sort of decide a certain configuration, on those sites, there will be forces. Atoms move a little bit away from that site, so that their force is 0. So what we have is here the sort of physical embodiment of the degrees of freedom. The energy is expanded up to the second order.

But then we also apply an equilibrium condition. That is, the force on the ions for a given configuration sigma must be 0. So the derivative, the first order derivative of the energy with displacement, that is the force, must be 0. And if you apply that at this and you work the algebra, you really obtain, for a given set of configurational variable, what are the displacements away from equilibrium that are induced by this configurational variable.

And of course, they are mediated by this response function of our system, the force on an atom induced by our configurational variable, and really, the second-order derivative, the stiffness of our potential energy surface with respect to atomic displacements. And so we use this equilibrium condition. That is, for a given configuration, the displacement are really given.

And we take it, and we stick it back inside. So we substitute, for the green u of R, this expression here. Those are all sort of tensorial quantities which I sort of removed the sums.

But so we are back to our spin Ising Hamiltonian in which the only thing that is relevant is the configurational variables, because at the end, that's the only thing that can happen. We can have a certain atom in a certain place, but we take into account the fact that the ions sort of go back, move to equilibrium. And we take that into account by a renormalization of the interaction constant j by, if you want, the displacement term or what we would call the elastic terms.

OK, so we have all the tools to calculate this regional j, this that would be nothing else than the dynamical force constant matrix that you diagonalize in your final calculation or this mixed term that in the literature are called Kanzaki forces. We do this on the virtual crystal. And then we really have our spin Hamiltonian, including the energetics effects of atomic displacements.

And so we have really the Hamiltonian of our system, and we can start doing the thermodynamics of it, with a small caveat that what we have done up to now is really this step. We have been able to calculate, if you want, the energy of a certain configuration, starting from a virtual crystal. But we haven't really-- and this is a subtlety that I won't go into the detail-- we haven't really taken into account that our sort of end original products have actually different lattice parameters.

So in itself, at the end, if we need to understand phased ability, we need to understand what is the lower Gibbs free energy for our problem, that is, if entropy drives this configuration to be stable, or, if you want, if enthalpy and elastic interaction drive segregation in two compounds with different lattice parameters to be the stable state. And of course, the closer the lattice parameters are, the more irrelevant this term is going to be and the more relevant term in which the disorder atoms will be. And temperature is always going to favor this.

But this might become very expensive if the original compounds have very different lattice parameters, because we really need to stretch things to start having the possibility of mixing them together. And all sort of-- the phase diagonal stability of this system is really driven by these two competing sort of forces, the temperature that want to mix up things around but chemistry that, in order to mix things around, you need to have atoms that are chemically very similar. And you need to have lattice parameters that are fairly similar, so that you don't build up a lot of strain.

And so, again, without going into detail, it is actually very easy to calculate some of those elastic terms. And so when we do that, what we really have is a way to calculate how much energy it takes to stretch things so they have the same sort of unit cell and then how much-- and that was the tricky part-- it cost to mix things around, it cost to have a certain configuration. And the very nice thing is that now, we can sort of test our perturbation theory approach and compare it with full blown density functional theory calculations for cells that are small enough so that you can actually do a density functional theory calculation.

These are older results. Now we could do much larger cells. But the concept stands.

We can do a full self-consistent calculation, say, for a super lattice of gallium planes and indium planes in the 111 direction. And so we can calculate what is the lattice parameter of this structure and what is actually the sort of Gibbs free energy of mixing for this system. And we can do that with a full explicit calculation, and we can do it with a set of perturbation approach.

And you see, we can compare the equilibrium lattice parameters in both cases. And basically, the difference-- this is atomic units. This is PWSCF-- is basically negligible. And the energies are also very good.

This sort of error that the perturbation theory makes is actually here, just 1 millielectron volt per atom. And as you sort of move around, you see that there are lot of cases in which the error is larger, 4 millielectron volt. And 4 millielectron volt is very small, but it's still 10% of this mixing free energy. So it's sort of somehow relevant.

But we, I would say, still do very well. And calculating this energy of mixing is absolutely inexpensive. It takes a nanosecond on your computer. You just basically sum with the appropriate configurational variables the term with your interaction constants. This requires a full self-consistent calculation.

So with this approach, you can calculate the energies of millions of structures that are very large, while with this, you can only calculate tens of structures that are reasonably small. So it's very appropriate to use the linear response term to then do-- to use one of these statistical sampling methods that you have seen with Professor [? Seder. ?] And again, I won't go into detail, because you have seen them partially, and you'll see them over and over again.

But just to remind you sort of the concepts, if you want to understand what is the stability of a certain structure, what you really need to compare are their Gibbs free energy. So you want to operate into the thermodynamic ensemble in which you have constant pressure and constant temperature, because that's really what your experiment is. You're at a certain temperature, and you let the system relax.

And as usual, you can't calculate-- explicitly, there is not a formula that gives you free energy. It's a logarithm of a partitioned function. And so there are sort of statistical ways by which that can be calculated. And one of the most powerful approaches is that given by thermodynamic integration.

That really is based on the fact that while we can't calculate free energy, we can actually calculate the derivative of free energy. So the derivative of the free energy with a certain parameter, like the derivative of the free energy, the Gibbs free energy with respect to the number of particles, is actually going to be given by the ensemble average-- that's when we use the sort of brackets, the bracket sort of symbol-- by the ensemble average of the derivative of the energy with respect to the number of particles. So that allows us to calculate derivatives of free energy in the appropriate thermodynamic ensemble.

And so what we can do is calculate free energy differences by integrating the derivative of free energy from state A and state B. And obviously, this integral will be done by step-wise calculating these derivatives at different points. But these derivatives, that in the case of the Gibbs free energy with respect to the number of particles, nothing less than the chemical potential as a function of the concentration, needs to be evaluated for every concentration. Say, in going from concentration A to concentration B, what we really want to figure out is what is the free energy difference between these two concentration. So each of these points in this integral requires an entire statistical sampling and then an entire statistical simulation.

And so this is what you would actually do. You choose a temperature, say, 1,000 kelvin, and you choose a pressure, say, 0 pressure or 1 atmosphere, that for all practical purposes is 0. And then under these conditions, you have your total energy expressed in configurational variables. And you decide what is your chemical potential, that is, what is sort of the additional cost of switching an atom from gallium to indium.

And then you start with your Monte Carlo sampling. You keep switching atoms from gallium to indium or from indium to gallium. Each of these switches will have a certain energy cost that is weighed with the Metropolis factor. And in this simulation, at a given pressure and at a given chemical potential, after of millions of sweeps and millions of trials in which you try to change a spin variable from plus 1 or minus 1, you do this over and over again, and you end up with a point there that tells you what is the concentration of, say, gallium 0.9 for a given chemical potential.

And you repeat this over and over again. That is, you sweep all your possible chemical potential. And for each choice of chemical potential, you get a certain average concentration. And so you do this over and over again.

And so you see, what you get is a curve that tells you what is the concentration as a function of the chemical potential. And you get one of these curves for each temperature at which you do the simulation. So you get concentration as a function of chemical potential.

And if we go back for a moment to the previous slide, this is really the inverse of what you need. You need the thermodynamic integration to calculate the free energy difference between two concentration, the integral of the average chemical potential for a given concentration. So the only thing that you have to do is, really, look at this curve, turn it around if you want to, so that you have x as a function of mu. And say, for the 1,000 kelvin case, you'll have something like this.

So you have sort of inverted-- your simulations have given you concentration as a function of chemical potential. You invert that, and you integrate it, and you get a free energy curve. And you see, things get interesting when you start being below a critical temperature.

So you see, we got 820 kelvin. And you really see that the curve is still continuous, but it sort of starts to have a slope in the middle that is going towards infinity. And if we go below, at 650 kelvin, what happens is that we have opened a miscibility gap. That is, there is a region of concentrations from 0.1 to 0.8 that never appear.

So at 650 kelvin, this system doesn't want to mix in the concentration regime between 0.8 and 0.1. And this is typical of model alloys. Miscibility gap has opened. The two things don't mix.

And if you think at sort of how your curve could look like, we could sort of somehow interpolate in some arbitrary way this curve and then invert it. And when you integrate the data-- it seems that the computer has gone mad. And let's see. OK, we'll keep the green forever. It's now in the realm of the computer.

But this is good, because when you invert this and you integrate, what you see is that for your free energy as a function of the concentration-- is that what you would really obtain is real free energies only from 0 to 0.1 or from between 1 and something like 0.8, 0.85. And in the middle, you have this famous miscibility gap in which, really, the free energy of the solid solution is higher than the free energy of the n compounds. And so you have the max error and the tangent. And so what you have is the composition of your system.

And so you can calculate your phase diagram. And at the end, for this system, it would really look like this. Below 820 kelvin, you open a miscibility gap in which, really, the two systems don't want to mix together. And so in many ways, semiconductor alloys are really model solutions. And below the critical temperature, they start to segregate.

And of course, you can do a lot of other intriguing things. That is, in your simulation, you figure out, say, what is the distributions of, say, gallium-phosphorus distance or indium-phosphorus distances at any given concentrations. And what you actually discover is that there is a sort of very bimodal distribution, that there are short distances with the phosphorus, and there are long distance with the phosphorus-- that are, of course, due to the fact that these are gallium-phosphorus distances and these are indium-phosphorus distances. But they sort of depends with the concentration.

But those are things that can be measured. And so in this sort of graph, I've put an example of what are the numbers that we get from the Monte Carlo simulations on the Ising Hamiltonian and what are actually the EXAES results that people measure. And you see, there is really a very good agreement. Say, these are the short gallium phosphide distances as a function of concentration between what the experiment measure and what your thermodynamical DFT is actually predicting.

And I think I'll skip this. There were some other details about epitaxial system. But otherwise, I'm worried I won't be able to do the second part.

So this was the first example. Again, what is important there is that we have a physical problem. We have identified what are the important degrees of freedom, and we have calculated those degrees of freedom from calculations that are actually manageable. But those degrees of freedom in those calculation have given us the parameter of a Hamiltonian that, in this case, was a Ising-like long range Hamiltonian in which we can do sort of much more complex statistical mechanics analysis. And so we can actually obtain phase diagrams.

That was one example. The other example that I want to show you here is instead how we sort of repeat the same process in order instead to calculate Hamiltonians that have a linear scaling cost. That is, that allow us to calculate the electronic structure and energies of systems that are really large, that are sort of thousands or tens of thousands of atoms.

And you know, this is because this comes from some research example. We were interested in sort of studying what is the effect on a carbon nanotube of organic functionalization. That is, what happens if we take a ligand-- this would be-- it's a benzene ring. So again, let's kind of call it phenyl moiety. And this sort of ring is attached with a covalent bond to an infinitely longer carbon nanotube.

And people have started doing this. And of course, when it happens, what you have is a nanotube that has a random array of these funny rings. So again, we have the thermodynamical problem of finding out what is the effect of this configuration of this order on some of the intriguing properties of the carbon nanotube. Like, there are sort of ballistic electrical conductivity, that is electrical conductivity, that is independent of length, so very different from an ohmic conductor-- or what is their thermal conductivity.

But again, we really don't have the computational resources, and no one has the computational resources to study, with density functional theory, a system with tens of thousands of atoms. And you'll see this in sort of the last part of the class, really. It's not just a matter of waiting a few more years. I mean, the scaling cost kills us in many ways. So we really need to use sort of ingenuity.

And what we are going to use here is this idea. We want to break away from this orthonormalization constraint that is really what drives the cubic scaling of a system. That is, what we want to do is using a system as sort of picture in which we really look at Kohn-Sham orbitals that are localized, so we don't have to worry about orthonormality with orbitals that are far away. Because if two orbitals are far away, and both of them are localized, their overlap is going to be 0. And so they are orthogonal by default, without us having to sort of impose that.

And the way we are going to do this is, again, a mapping from the full density functional theory calculation. So we start-- I'm giving you here the example of silicon. If you look at the band structure of silicon, again, what you have are very localized core orbitals down there that don't really have any effect on the chemistry. They corresponds to flat bands.

But what is interesting is what the valence electrons do, and this is sort of the very celebrated band structure of silicon. This in particular is the gamma point, so at 0, 0, 0. And so because silicon has two atoms per unit cell, four valence electrons per atom, we have eight valence electrons. Spin the generator. So we have bands that contain two electrons each. Sorry.

So what we have is four valence bands. And you see, 1, 2, 3, 4. As we move around the Brillouin zone, they can become the generator. And those that are studying semiconductor physics, they will know that this is the top of the valence band, and we have heavy holes and light holes depending on the curvature up here.

But so what is the electronic structure of our problem? Well, remember, the band index and the quasimomentum k across the Brillouin zone are really the quantum numbers for the crystal, for my problem. If I study a molecule, and I have a molecule with eight electrons, I will have only four-- in the case of spin degeneracy-- eigenstates that are occupied.

If, instead of studying a molecule with eight electrons, I study a crystal with eight electrons per unit cell, instead of having four states, I have four bands. And the energy of each eigenstate, thanks to Bloch theorem, can be cataloged depending on this band index and depending on the quasimomentum. And the quasimomentum really sort of tells me how the phase of things changes when we go from one unit cell to the other.

The same concept takes place also not in-- for the electronic eigenstates but for the normal modes, for the phonon modes of a crystal. We classify them according to wavelength and according to band index. And so your infinite number of electrons in the crystal can all sort of be summarized in this picture, in which we have a continuum of states. So if you want, we have mapped the fact that we had eight electrons repeated infinitely into a problem in which we have eight electrons, that is four bands, with a sort of continuum variation, depending on which wavelength we are looking at. And so this picture sort of represents all the possible occupied states in your system that are infinite but are classified in a group of four with a continuum index k.

Now, the interesting things about quantum mechanics is that you can change your representation and still look at the same problem. It's the same thing in classical mechanics. If you are sort of studying something in three dimension, you can rotate your coordinate system. And so your equation of motion will change. Your trajectory will change. But really, you are just moving everything accordingly, and you are looking at the same problem in a different coordinate system.

That is even more true in quantum mechanics, in which you can perform a unitary transformation between all your eigenstates and still look at the same problem. Usually, what we sort of choose to be is into a representation in which we are looking at eigenstates of the Hamiltonian. So each of these states here is exactly the characteristic of being an eigenstate of the Hamiltonian. You apply the Hamiltonian to that state. What you obtain is a constant that is the energy eigenvalue times the state itself.

But we can mix this around, and we can still look at the same problem, just in a different representation. Exactly as in classical mechanics, we can do an orthogonal transformation of our coordinate system and look at the same problem. So what we can do is actually do a unitary transformation, a unitary rotation, of our orbitals and still look at the same problem.

And again, without really going into the details of what it is and what quantum mechanics is, it's something that would look like this. That is, if we start from our eigenstates, at each k point-- that is, at each point in the Brillouin zone-- if we have those four bands, we can mix them together with a unitary transformation. But not only that, we can sort of mix those mixed bands with all the other bands at different k points with the appropriate unitary factor.

And so what we can obtain is a new set of orbitals that are, again, infinite and that really represent the same quantum mechanical problem, just sort of turned it around. And our quantum numbers that before, where the band index and the quasimomentum become some new quantum numbers, that will be-- again, one of them will be related to the bands. So for silicon, we'll go from 1 to 4. And the other one will count all the sort of infinite number of electrons that we have when we go from one isolated unit cell to one unit cell periodically repeated.

And before, what we had was a quasicontinuum, basically, of wavelengths. And in this new representation, what we have is a quasicontinuum of localization vectors. So we go from extended Bloch orbitals that are eigenstates of the Hamiltonian to localized orbitals in which the electron is sitting in each of the different unit cells.

And really, this is an arbitrary transformation. And these matrices are sort of our parameters, the sort of quantum mechanical equivalent of the three Euler angles that rotate the Cartesian system from one representation to another. And so we can choose, say, this set of transformation, so that what we obtain is the same quantum mechanical picture but in a representation that is as localized as possible.

And how we do that sort of requires a lot of algebra. But this is the sort of very intriguing result that takes place when we actually make it in practice. We take-- again, all these states and each of them is delocalized. It's Bloch orbital energy eigenstates that is extended over the whole crystal.

We mix all these things together with a criterion that those mixing transformation are chosen so that the new states are as localized as possible. And what we obtain at the end is fairly beautiful. If we look at, say, something like silicon, again, sort of a perennial diamond structure, one silicon atom fourfold coordinated with the other silicon atoms-- and it's a bit difficult to see, but what we would have is that one silicon atom, in red, is connected to four silicon atoms. And this blue over blue silicon atom is connected to other four silicon atoms.

And the localized orbital that you obtain is really what the chemist would have told us probably in 1930. It's really a covalent bond between the two orbitals. So what we have is that these two atoms here are connected by an orbital that has, really, this pile up of covalent charge between those two atoms. And there is a certain amount of sort of back-bonding between each of those two atoms and the remaining three neighbors.

So this is how the localized orbital looks like. And that's been just obtained by a unitary transformation. But when we have this player here, it becomes very easy now to construct the electronic structure of a system that might have some amount of configurational disorder. Because, if instead of having somewhere, say, silicon-silicon in an infinite silicon matrix, we have maybe gallium and arsenic for a moment locally.

Well, maybe what we can do, and I'll show you how we'll do it, maybe we can put there the exact orbital corresponding to the gallium arsenic bond. And if you look at it, it's actually going to be-- again, because things are similar, it's going to be very similar to the silicon orbital. The gallium arsenic covalent bonds is just, because of the difference of electronic activity, sort of tilted more towards the electronegative arsenic in this problem.

Or say, again, to give you the feeling, it's nothing less than the sort of standard chemical picture. If we have a system that is still silicon, but it's amorphous silicon-- and I've got here atoms around. But amorphous silicon is something in which a tetrahedral short range holder is still preserved on average but long range holder is lost.

Well, atoms that are tetrahedrally coordinated won't have a perfect 109 degrees angles, but their bonds still look very much like those. And sometimes, you have atoms that are five-fold coordinated. And so they have their own bond.

And if you want, these are really what we call the LEGO bricks of electronic structure. That is, in a very complex system, the electronic structure of the system is going to be made up locally by the appropriate orbitals that corresponds to the local environment. And the nature and physics is inherently local, so there is a large degree of transferability.

And again, to give you a sort of more feeling for this, suppose that instead of looking at a crystal, we look at a molecule. We look at something like a benzene ring. And you go to a representation in which states, instead of being eigenstates of the Hamiltonian, are localized states.

What you see are electronic states that are really carbon-hydrogen single bonds, or what you see are carbon-carbon single bonds, or what you see are carbon-carbon double bonds, in which you have both this orbital and this reflection across the symmetry plane of the benzene molecule. So really, what you are seeing is very intuitive chemical picture. But the way it's been obtained is actually from an exact transformation of our quantum mechanical problem, solving density functional theory, that is, with a certain degree of accuracy.

And we'll see now how we apply this to the case of carbon nanotubes, because what we really want to do is study chemical disorder on these carbon nanotubes. And carbon nanotube is nothing else than a sheet of graphite-- that is what is called the graphene sheet-- that has been rolled over together to create a tube. So if you take this hexagonal lattice, there are sort of discrete infinite way of rolling things together.

And the way the carbon nanotube community has agreed to sort of use a nomenclature is by defining what is called a chiral vector, that, in this case here, would be this CH vector. So if, on a graphene sheet, you define a vector, that is really just an integer number of times the two a1 and a2 primitive lattice vectors-- that is, once you have defined this green CH, you have defined your nanotube, by saying that you are cutting the graphene sheet perpendicular to this vector and rolling what you have cut, so that the top of the vector touches the bottom. And so you have sort of a perfect cylinder.

And then there is a nomenclature. And you can sort of-- a generic nanotube will be called armchair, zig-zag, or chiral nanotube, depending on how you have rolled things around. And so if your chiral vector is an n number of a1s and an n number of a2, equal, your chiral vector will be this blue vector here.

And so what you are going to do is, you are going to roll the nanotube so that perpendicular to the direction of the tube, you can clearly see this sort of zig-zag pattern. And so an armchair-- sorry, a zig-zag-- oh, sorry. A zig-zag nanotube is the one in which you have a chiral vector that is made by an integer number of times a1 and zero number a2 or vice versa. So in the form n, 0 or 0, m. And so you see, this blue vector for the zig-zag is going to be 1, 2, 3, 4, 5 times a1. So this will be the zig-zag nanotube, and you see this zig-zag pattern.

If, instead, you choose n for a1 equal to the n for a2, what you have is a vector that goes in this direction, like the red vector here. And so what you are doing is actually sort of cutting the nanotube like this and rolling it together. And again, now perpendicular to the axis of the nanotube, you see this pattern. That is called an armchair pattern.

And so an armchair nanotube is identified by these two quantum numbers. And then identical zig-zag is n, 0 or 0, n. And everything else has a sort of more ordered chirality and is called, generically speaking, a chiral nanotube.

OK, so this is how they look like. You see, a zig-zag, very long in this direction. We cut it, we see the zig-zag pattern. An armchair, we cut, and we see that sort of hexagonal pattern that, for some reason, is called an armchair.

And now the interesting bit comes when we look at the band structure of this system. We could have started, say, from the band structure of graphite or graphene, in which, again, the Brillouin zone-- the extended Brillouin zone is this hexagon. And what happens is a bit difficult to visualize. But actually, what happens is that there are only six points-- so the black dots there-- at which they are unoccupied states touch the occupied states.

So usually-- and we have seen it over and over again-- in something like a metal, we will have maybe a band like this. And then we'll have Fermi energy that cuts the occupied states from the unoccupied states. In graphite, the band dispersion is much more exotic.

It actually looks like this. In certain points of the Brillouin zone, it's actually sort of straight lines. You'll see it in a moment.

And the Fermi surface is exactly here. So there is only-- the Fermi energy is here. The Fermi surface is on a point. And seen from above, if you want, these are where these combs are that I've sort of cut there in a profile.

What happens when you go from graphene to a nanotube? Well, instead of having a system that is infinitely B-dimensional, you get a system in which, by rolling one of the dimensions, what you do actually is you quantize your quantum numbers in one direction. So you still have a Brillouin zone that is a continuum in the direction of the tube axis. But because you have rolled it, maybe you have only 10 hexagons instead of an infinite number of hexagons in the direction perpendicular to the axis.

Now, what you are effectively doing is really looking at the band structure only along a certain number of cuts of this Brillouin zone. And so depending on how you roll your system, what you can have is that this subset of lines that corresponds to having a continuum index in one direction become a discrete index, because in one direction, we go from infinite hexagon to a finite number of hexagons. And so depending on how many and what is the orientation of the cut, we can have that these lines cross or do not cross these points.

And if they cross those points, we'll have a metallic nanotube. If they don't, we'll have a semiconducting nanotube. So just a very slight change in orientation in the chirality vector might allow us to miss this point and make something that was metallic into a semiconductor.

And the rule to remember this is trivial. If the difference between the indices is 0 or a multiple of three, your nanotube is going to be metallic. Otherwise, it's semiconducting.

So say, all armchair nanotubes are metallic, because m is equal to n, so the difference is 0. So they are metallic. Zig-zag nanotubes, well, only 1 out of 3 is going to be metallic.

6, 0 is going to be metallic. 7, 0, 8, 0 are semiconducting. 9, 0 is metallic, and so on and so forth.

And this sort of only breaks down when we have very narrow nanotubes in which this sort of model picture of what is called band folding starts to break apart, because the way we have thought of the nanotube is as being identical to a perfect sheet of graphene that has been quantized in one direction. But really, there is curvature. And if the nanotube is narrow, there is a lot of curvature. And so that sort of might change things around a little bit.

But so here is an example. You see, if you have a metallic nanotube, like in this case, when you quantize it in one direction, you are rolling it, you actually end up with lines that go through the black points. And so what you have is that your band structure there sort of sees these cones of the graphite system.

If you have, say, in this case a semiconducting a0 system, when you quantize anything in that perpendicular direction, your lines actually do not go through the black points. And so you never encounter a corner, and you never find that gap. And again, a very slight change in orientation can give you a very variable gap and can bring you to be metallic very quickly.

And so we start really going again into the details. We really need to take all our occupied states-- that is, the one below here that I'm not showing, everything below the Fermi energy that is here for this metallic nanotube-- and we actually want to take a little bit of the states above, because those are also important. And with all these states, the red, the blue, and the black-- so all of the occupied states and a bit of the unoccupied one-- we want to sort of mix them together and find what are the localized orbitals coming from this.

And I won't go into the details on how we sort of extracted those when you have some of the discontinuity that comes about in a metal. But so basically, what's happening is that we start from our electronic structure calculation, maybe of the carbon nanotube with a functionalized ligand. And in a new quantum mechanical representation, what we have obtained are sort of the building blocks of the pristine nanotube, and the building blocks of the ligand, and things like what is exactly the orbital that connects a carbon atom in the ligand to a carbon atom in the sidewall, with sort of full first principle precision.

And so with these blocks now, we can actually construct not only the band structure and the electronic structure of this system periodically repeated, but also a system in which the ligand is first here, here, and then here, with all the chemical disorder that we want. And because the orbitals do not overlap with each other, the computational cost of doing this is very, very reduced. Because basically, in our new representation, orbitals that are far away are orthogonal with each other, when we actually represent the Hamiltonian-- remember, in a lot of your PWSCF calculation, you have represented the Hamiltonian in a plane wave basis set. And so in a plane wave basis set, your Hamiltonian has the diagonally dominant part, that is g squared over 2. But then it also terms all over the place that corresponds to the Fourier transform of the potential.

When we actually represent the Hamiltonian on this basis of localized orbitals, what you have is that this is truly diagonally dominant. And as soon as we consider orbitals that are far away from each other, there is no Hamiltonian matrix between them. If we calculate the matrix element between this, the Hamiltonian, and this green one, it's going to be 0, because they really do not overlap. So this is what is called a sparse diagonally dominant Hamiltonian. And diagonalizing this is something that actually doesn't require cubic scaling but only requires linear scaling. So one can really sort of study very large systems.

And to show you an enlargement of a cutout from a perfect nanotube, you see, this is what you get. You get orbitals that look either as beautiful sort of bonding combination of sp2 orbitals-- if you think, we know from chemistry that the carbon is sp2 hybridized. So it means that the graphitic backbone is given by a bonding combination of an sp2 orbital sitting here and an sp2 sitting here.

And the same happens here. And then there is a set of pz orbitals sitting on each tube. And when those mix together, they give rise to the metallic graphitic-like bonds. And again, this looks a lot like a pz orbital, but it is nowhere similar to the pz orbital of a carbon atom. It's the pz orbital specific to a carbon nanotube.

So with this, we can-- as previously in the case of the alloys, we can first validate our system. That is, we can start constructing the electronic structure of something that we can also fully calculate from first principle. And we can compare them and these continuous lines from the models and dots from the full calculation.

And you see, the model works perfectly and doesn't describe the higher energy bonds that have not been included in our sort of orbital construction procedure. But for all the bonds that we wanted to consider, this is perfect. And once we sort of feel confident that this works, we can actually use it to calculate the-- let me actually go farther-- we can actually use it to calculate the electronic structure of systems that are too large to attack directly, something like this.

This has just a conducting region of 3,000 atoms and a very long sort of semi-infinite leads on the side. And we can figure out a lot of what happens, say, to the ballistic conductivity of this system. And I'm not really going into the details.

I wanted to go before concluding to a previous slide here, in which I was actually showing how, in practice, the matrix element between a localized orbital and another localized orbital changes as a function of the distance, say, as we sit on a carbon atom, and we look at the overlap between the pz orbital here and we move away from the nanotube. And you see, in a semi-logarithmic scale, we have very good exponential decay. Certain point goes into the numerical noise.

And we can really look at-- validate what are the effects of our truncation by cutting shorter and shorter Hamiltonians, by neglecting a large number of terms that are far away. And you see, there is almost an invisible effect on the band structure. So that really sort of tells us how much localized things are, how much they overlap, and if you want, tells us how long range or short range chemistry is. That is very important when we consider one of this configuration.

So with this, I'll conclude. Again, the two important messages that you want to remember from this class is how you take a physical problem, and you can use density functional calculation to find out what are the important parameters and construct a model for your energy. The first case where configurational variables inducing relaxation and gives us Ising-like Hamiltonian, in this case, we have localized orbitals. And once we determine those, from a small calculation or from a fragment, we can actually calculate the very large system.