

PROFESSOR: So what we'll do today is really switching to the computational description and characterization of thermodynamic properties. As we have seen in a lot of the previous lectures that we have been focusing on finding out good energy models, that is being able to calculate, what is the energy of a system given the coordinates of its ions?

And now what we want to understand, really, is, what is the evolution of the system in the presence of a finite temperature? Temperature really means that there is an average kinetic energy available for all the atoms and the molecules in your system. That means that there is really a microscopic dynamics. And we want to follow that.

One of the most important things-- and this goes back to, really, the first lecture that you have seen in this class for the quantum part-- is that there is the fundamental de Broglie relation that gives us an estimate of what is the quantum wavelength of any kind of object. Remember, we have said that the wavelength times the momentum should be equal to the Planck constant.

And that really means that for an object like an electron, the de Broglie wavelength that comes out from that relation is comparable to the distance between the atoms. So if the wavelength of an electron is of the order of the angstrom, 10^{-10} meters, it means that the wavelike properties of the electrons needs to be studied. And the electron will behave as a wave. And that's why we use the Schrodinger equation to describe electrons.

The wavelength of the nuclei is much, much smaller. That means that nuclei really didn't-- we don't really need to take into account to a good approximation the quantum nature of the nuclei. The wavelength is so small that they can be treated practically, correctly, as classical particles.

And that's why, actually, nuclei, in a way, still follow the rules of classical dynamics. That is, they evolve in time following Newton's equation of motion. Again, this is an approximation that involves treating the nuclei as classical particles and not as quantum particles. But especially at high temperature, this is a very good approximation.

And so a lot of approaches in computational and material science have to do with solving Newton's equation of motion for a system of interacting nuclei and evolving those nuclei in time to study the thermodynamic properties of the system. And so that's why I actually started with a reminder of Newton's equation of motion. That is really an ordinary second-order differential equation.

That is, when you have a particle of mass m in a force field that is this F of r , well, we can calculate its trajectory-- that is, its position-- as a function of time and its velocity as a function of time by integrating this differential equation. And in the classical case that is always presented, say, is the evolution of a particle in a constant force field, like a particle under the gravitational law, that, to a large extent, is constant if you don't move very much.

And because, in that case, the force is really a constant, the integration of this becomes trivial. It's really a second derivative that needs to be equal a constant. And integrating that gives us a parabola.

But the reason why I'm focusing on this is to remind you of the initial conditions in our differential equations. That is, if we have a second-order differential equation, the general solution-- let's say in this case for this force field would be a parabola-- is uniquely identified once we specify the position at a certain instant and the velocity at a certain instant.

That is, among all the possible parabolas in the world, you selected the one that your projectile is going to follow once you have the position and the velocity at a certain instant. And that's why actually dynamics and molecular dynamics and other physical properties basically depend essentially on the position of the particle in your system and the velocity of the particle in your system.

That's why there is no physical observable that depends, say, on the fifth derivative with respect to time. It's really because the dynamics of that system follows Newton's equation of motion.

OK, with this, what we really need to deal in general is with a system of interacting particles. And for a moment you could actually think that this i , interacting particles, could be the planets orbiting around the sun. So each of the planet i will have its own trajectory. Our r of t will have its own mass.

And they all live in a force field. That's really where the many-body complexity, again, comes out. So many-body complexity, that is much easier to deal with than the many-body complexity of the electron problem because, again, now we describe each object just with a set of six variables, so its 3-space coordinate and its 3-momentum coordinate.

But still, each and every object interacts with each other. When you look at the planets orbiting around the sun, you, at first approximation, could only consider the gravitational attraction to the sun. But in reality, also the planets interact with each other.

And so, again, even if this is much simpler than the Schrodinger equation, as soon as you start to have more than two bodies interacting, it really becomes too complex to solve analytically. And so all the molecular dynamic techniques have, as their only object, basically, the integration in time of this equation of motion in an efficient form.

There are a couple of comments that one wants to make. In general, we deal with conservative fields. That is, the force depends on position only. That also means that the work that you make in going from one point to another point, the integral of the force times the displacement is going to be independent from the trajectory. So the energy, if you want, is just a function of the position.

The other thing that descends immediately from this equation of motion is that the total energy of the system is conserved. And that's very easy to prove. And I'll do it in a moment.

And that's what's often called a microcanonical evolution. When you read about microcanonical ensembles, it means that you are actually under the condition in which the total energy of your system is conserved. And again, what is the total energy of your system? Well, if you have one particle, it's just, trivially, the kinetic energy $\frac{1}{2} m v^2$ plus the potential energy V of r of that particle, and just doing everything in one dimension.

And you ask, why is this going to be conserved? Well, you can just prove it by taking the time derivative of this quantity, OK? So we have taken kinetic energy plus potential energy. We look at the time derivative. And this is just simple algebra.

What you get is $\frac{1}{2} m$ times 2 times v times the derivative with respect to time, that is the acceleration, plus the derivative of the potential with respect to time. And this can just be written as mass times velocity times acceleration plus-- and we can rewrite the derivative of the potential with time as the derivative of the potential with respect to position times the derivative of the position with respect to time.

And now you see where we are getting to, because this term here is equal to minus the force, while these terms here is equal to the velocity. And so it goes away with here. And so what we have is that m , the derivative of the total energy with respect to time, is equal to the mass times acceleration minus the force. That is 0 because the particle follows Newton's equation of motion.

So because of Newton's equation of motion and because the particle follows exactly that evolution, we have that the sum of the kinetic energy plus the potential energy needs to be conserved during the evolution. And that's actually the fundamental sanity check that you will always have and you can always do during a molecular dynamics simulation. That is, you can check that your integration is correct by making sure that the total energy of your system is conserved.

OK, so what do we do in a practical molecular dynamic simulation? Well, the idea is that, as in the motion of the planets, we want to integrate the equation of motions with this caveat that, again, each particle interacts with each other particle. And so it's really a many-body problem, even if it is a classical many-body problem. And so it's really easier to integrate than a many-body Schrodinger equation.

And in general, the way particles interact with each other can be described by empirical force fields or, in the best-case scenario, can be described using quantum mechanics. That is the [? functional ?] theory [INAUDIBLE].

And so what you will have to do in that case is that for every configuration of your ionic system, you need to solve an entire total energy self-consistent problem. And that's why actually the first principle in molecular dynamics becomes a very, very expensive challenge. And we'll see in one of the later classes how this is solved.

There are a lot of systems that can be described accurately, even with very simple potential. Even more than in the case of the total energy, certain thermodynamic transitions, like the transition from a solid to a liquid, can be thought of as broadly independent of the details of the interacting potentials.

So even studying a system with a Lennard-Jones potential that, after all, was developed especially to do molecular dynamic simulations of rare gas atoms, or even just studying a system that is made by hard spheres that don't really interact when they travel in space, apart when they really hit each other-- so what is called a contact interaction that is 0 all the time apart from the instant in which the two spheres collide-- can sometimes give you a very accurate qualitative description of the thermodynamic transition in a problem.

And molecular dynamics actually was first introduced really to study systems, like liquids, for which it's very difficult to develop analytic techniques. You see, if you want to study a solid, you can actually choose a potential. And if the potential is simple enough, you can analytically calculate the potential energy of a periodic repetition of all your atoms.

But if you study liquid, disorder is really an essential part of the state of the system. And so the entropic contribution to the free energy of the system is fundamental. And really there are no reasonable analytic ways to address a lot of the questions that you have. And that's why people really moved to solving on the computer the Newton equation of motion.

And this is a short history, really, of molecular dynamics, reminding you what is the history of computers, really. In 1952, one of the first computers, MANIAC, was operational in Los Alamos. And that's when Metropolis, Nicholas Metropolis, in collaboration with actually a husband and wife couple, Rosenbluth and Rosenbluth and Teller and Teller, developed what is called the Monte Carlo method.

That was a very powerful method to address exactly this entropic free energy problem. And that you'll see in one of the later classes. But really, 1956 sees the first molecular dynamic simulation by Alder and Wainwright.

And as I said, in looking at the dynamics of hard spheres, these are the milestone papers. Vineyard is studying radiation damage in copper, what happens when you have a copper atom that is really put in motion by a very energetic radiation kicking it and starting running around.

And really this is the classic paper of Aneesur Rahman studying, with a Lennard-Jones potential, the dynamics of liquid argon. Again, a Lennard-Jones potential is actually very accurate to describe rare gases that ultimately repel each other at very close distance and, at large distances, attract each other with this weak dipole-dipole interaction that goes as $1/r^6$.

And then really the first ab-initio molecular dynamics and the first theory for ab-initio molecular dynamics was developed in the mid '80s by Car and Parrinello. OK, now, there is a universal way to think of these dynamics of your classical system composed of N particles. And as we have said, since they follow Newton's equation of motion, what we really need to track, the only quantities that we need to track, is their position as a function of time and their velocity as a function of time.

And generally speaking, the kinetic energy is just a function of the velocities. And the potential energy is just a function of the positions. So if you have N particles, you really have to follow in time $6N$ variables.

And if we define for a moment a six-dimensional space, what we really have is that an instantaneous state of your dynamical system, that is all its position and all its velocity, can be represented by a single point in this $6N$ dimension of space. So if you have got just a hydrogen molecule with two nuclei, what you really need is a point in six dimensions.

If you have a water molecule, you need a point in nine dimensions. And the evolution-- sorry. If you have a water molecule, you need a point in 18 dimensions. And if you have a hydrogen molecule, you need a point in 12 dimensions.

And the dynamical evolution of your system is described exactly by the evolution of this point. So we often think at this phase space that we'll have, for an N particle, $6N$ dimension, of which-- let's try to draw a bit dimension-- $3N$ of them, those represented in green, will really represent the positions. And the other $3N$ will represent velocities.

OK. And obviously this is all orthogonal dimension. And your point will be somewhere. The certain instant in time T will have a set of velocities, will have a set of positions. And the Newton equation of motion will make it evolve in this phase space.

So really your molecular dynamic integration consists in following in time this evolution. And that evolution is, in principle, analytically and uniquely defined once you know the position and the velocities at every instant. And-- well, we'll see that later.

OK, so what do we do once we have a computational algorithm that allows us to evolve this point in phase space? And I sort of summarized here the three main goals that could follow from a molecular dynamic simulation in the order that we'll see them.

Really, one of the most common ones is to use molecular dynamic simulations to calculate thermodynamic properties, that is to calculate what we call ensemble averages. Suppose that you have a system like a solid. And you have calculated, say, its lattice parameter, like you're doing in lab 2 or lab 3.

And suppose that you start heating up this solid. What happens? Well, the atoms, in the classical sense, start oscillating more and more. And you know from experience that most solids start expanding when you warm them up.

So a molecular dynamic simulation could tell you how much a solid expands when you heat it up because basically you start evolving your system from 0 temperature. And then you increase its temperature. And really, temperature, remember, is just the average kinetic energy of the ions.

And as you increase the temperature, the ions oscillate more and more. And if you let your unit cell expand or contract in time following what is the stress overall of all the atoms in your unit cell, you see that your system increases the dimension of the unit cell and starts oscillating around a value that's really an average, depending on temperature.

And all these sort of statements can be actually made formally in a way that we'll see in a moment that is the language of thermodynamics. That is, if we have a finite temperature, we'll have a certain probability associated with every microscopic configuration.

OK, so if you are at 0 temperature, really what you have is that you have a probability 1 to be in the lowest energy state possible and probability 0 of being in any other energy state. If you start increasing your temperature, states that are energetically similar to the ground state but not really the ground state start to have a certain probability of being occupied.

And the more temperature you have, the more very costly configurations can be accessed by your system. And really, the average properties at a finite temperature need to be an average over all these possible states weighed with their appropriate thermodynamic weight, weighed with the probability of being in that state.

And so this is what an ensemble average does really. It tries to span all the possible states in a system with a probability that is proportional to their thermodynamical weight. And then for each state, you can calculate a certain property, like what is its average potential energy, what is its lattice parameter, and so on. And out of that, you can calculate the average thermodynamic property.

And we'll see, actually, an example in a moment. So we'll see a practical application of this. The other thing that you can do with molecular dynamic simulator is actually study the evolution in real time of your system. So you could actually study a chemical reaction.

What happens if you have, say, an explosive molecule that is decomposing, OK? And in particular, whenever you have a chemical reaction that tends to be a bond-breaking and bond-forming reaction, you'll probably need to use a quantum mechanical approach.

But in principle, you can set up your system in an initial configuration and see what happens as you evolve in time. You could look at a catalytic reaction. Say, maybe you are interested in studying fuel cells. And you want to see how a hydrogen molecule decomposes when it arrives on a platinum surface.

Well, you can take your hydrogen molecule and project it delicately towards the platinum surface. And follow, with your molecular dynamic techniques, the dynamics of all these nuclei as the molecule chemisorbs and dissociates in time.

And then really the last application of molecular dynamics is more of an optimization algorithm. There are problems in which it's very complex to find what is the lowest possible solution, what is the optimal solution.

That is, you could try to find out what is the schedule of all your flights in an airplane company. And obviously that's a complex optimization problem because you can't move one plane at a time and figure out what is the best possible solution because there might be a completely different choice of planes and itineraries that actually gives you a best performance overall on the net.

And so you can always take one optimization problem and express a cost function. That is depending on what is your interest. Minimize the amount of oil, let's say, used in running all these routes.

So when you have a cost function of a very complex problem, you have really an object that depends on many variables and that has a lot of possible minima. OK, so this would be, really, the problem of optimizing the air routes of an airline. What you have is a lot of possible variables. And there are a lot of reasonable solutions.

But you really want to find the one that is the best for you, something like this. And molecular dynamics can actually help to solve this problem because it's a problem that is absolutely analogous of having a system of interacting particles that has a certain potential energy surface and trying to explore that potential energy surface, moving around, until you find the lowest minimum.

You could think of this as a mountain range. And you have your skiers that are really your particles moving around. And you want to sort of let the skiers ski around as much as possible until they find the minimum energy state.

And so you can use a thermodynamic analogy. You heat up your system. That really means that every particle has a lot of kinetic energy. So having a lot of kinetic energy can easily overcome any kind of potential barrier.

And you have a lot of these particles moving around. And then you cool them down slowly. And as you cool them down, they'll start piling up in the lowest possible energy state.

And once you have reached, very slowly, 0 temperature, you look at where all your particles have ended up. And so you have at least a good sampling of the relevant minima in your potential energy surface. So actually molecular dynamics-- and this is called simulated annealing, in this context-- can actually be used as an optimization technique.

And we don't have to go as far as the case of an airline or what is called the traveling salesman problem. But often what we have is that we have a system that is fairly complicated to characterize.

Say we want to describe maybe an interface between a perfect silicon crystal and the SiO₂ substrate because we are in the electronic industry. Well, I mean, we have a problem because how do we construct the interface between a crystalline solid and an amorphous system?

I mean, we can try putting atoms in a reasonable position. But it's never going to work very well. And so what we can do is actually use molecular dynamic techniques to let the atoms evolve according to their interaction and to the Newton equation of motion and somehow find, by themselves, a slightly more favorable green minima instead of the initial configuration in which we have put them.

OK, so, in principle, if you had infinite computing power or infinite mathematical prowess, you would have really solved the problem of characterizing on a computer any material that you wanted. The reason why you can't really do that in general and indiscriminately has to do with four problems that arise in the approach that I've mentioned.

And I've mentioned them here, I guess, in order of importance. And the first and most dramatic problem is the one of time scales. That is, atoms move around and vibrate at time scales that are characteristic of atoms.

That is, say, in a molecule, it will take tens of femtoseconds or hundreds of femtoseconds to have a periodic oscillation. So these are femtoseconds. 10^{-15} seconds is the order of magnitude of the dynamics of the atoms at room temperature.

Now, a lot of interesting problems take place in times of seconds, minutes, or hours. A classical problem is how a protein folds itself in its native state. And you know, this is a process that can take seconds.

And so we have, really, the issue that all our calculation takes place on the timescales of 10^{-15} . And often we need to reach physical properties that are of the order of seconds. So we need to span a 15 order of magnitude.

And that means that if it takes a time x to perform one single molecular dynamic operation, we have a 10^{-15} cost to get to the reasonable time scale for a lot of macroscopic processes. And that's really a scale that we can't bridge.

If we do quantum mechanical simulations for a small system, we can easily reach the times of tens of picoseconds or hundreds of picoseconds. If we do classical molecular dynamic simulation, we can go maybe three, four, five orders of magnitude better. And so we can start accessing nanoseconds and maybe start getting closer towards the microsecond.

But there is still an enormous gap to reach. And this is really a difficult problem to overcome because we have dynamics taking place at different levels.

Say, think of a protein. We have the atomic dynamics of an atom vibrating around. And then we have dynamics of subunits maybe interacting with the different amino acids along the chain.

And then we have structural motifs that want to move around and then maybe the first coil-up in a certain configuration. But that's not really very good. And then after a while they sort of open up and recoil in a different configuration.

And so the more you look at your system on a larger and larger length scale, the more you discover that there are dynamics of more and more complex groups that are slower and slower. And integrating from the atomic motion up to the mesoscale and the macroscopic scale of these dynamics is really an impossible problem to do by brute force.

And I would say that it's still the conceptual problem for which less progress has been made. It's really the most difficult one. And you'll hear a lot of talks nowadays in science about multiscale modeling and about bridging timescales. And it's obviously a very important challenge. But I still think we are at the stage in which there are very little solutions.

Somehow related to the time-scale problem, there is the length scale problem. Suppose you want to study a system close to a phase transition. Suppose that you want to see a gas evaporate. You want to study a liquid-to-vapor transition.

What happens exactly at the temperature where the liquid and the vapor are in equilibrium? Suppose that you are studying water boiling. At normal condition, it boils at 373 Kelvin. And at that temperature, why there is a phase transition? Basically because the free energy of the liquid and the free energy of the gas are the same.

You go a little bit above the temperature, the free energy of the gas is lower. So the system wants to be all gas. You go a little bit below the boiling temperature, the free energy of the liquid is lower. So the system wants to be a liquid.

But as you get closer and closer to the temperature, these two free energies becomes comparable. That means that the cost of transforming from liquid to the gas and from gas to liquid becomes smaller and smaller. And it is 0 at the transition temperature.

That means, in practice, that since it doesn't cost anything to transform from one state to the other, you have lost all possible length scale because you can nucleate a bubble of gas of any size with 0 cost. So if you think of the transformation from a liquid to a vapor as the nucleation of bubble of gas exactly at the temperature of the transition, that cost has become 0. And we can nucleate bubbles of any size.

So the phase transition is a moment where we lose length scales. The fluctuation can have any size. And so any kind of finite simulation will break down because it's really only in the infinite system that we represent appropriately all these sizes popping up in our system.

And so that's why it becomes very difficult to study accurately, say, a boundary at a phase transition. We say that the length scale starts diverging. And that's why one needs to be careful.

This problem is not as serious as the problem of time scales. But it's still one that you want to consider very carefully and one you need to spend some time thinking, actually.

That is, often we use a periodic boundary condition to describe an extended system. And so the question that you always have to ask yourself is, how large should I make my unit system that is periodically repeated before it's really describing accurately the infinite system that is the goal of my simulation?

Suppose that you want to study water, OK? Or suppose that you want to study any liquid. In principle you want an infinite system. In principle you can deal only with a finite number of degrees of freedom.

So the two possibilities that you have, you could study a droplet of water, OK? But then obviously you encounter the problem that, when you study a droplet, you have a very large influence coming from the surface effect.

A droplet is a droplet because of the surface tension. And before reaching the thermodynamic limit of that droplet behaving as a bulk liquid, you need really to go to enormous sizes.

If instead you study, say, the same number of molecules in periodic boundary conditions, you eliminate the presence of the surface. And so you can reach the thermodynamic limit of the system behaving as a bulk much faster.

But still you need to always ask your question, is your simulation still large enough? And there is a size that is large enough. And that's the other very important concept. That is, whenever your simulation cell is so large that the dynamics of one water molecule doesn't affect the dynamics of a water molecule that is far away from yours as much as it is its periodic image, your system has become infinite.

That is, whenever you have a molecule in a liquid, you really dynamically interact not with every other molecule in the infinite liquid but only with a sphere of neighbors. And you need to understand how large is that sphere of neighbors because those and only those are the molecules that affect yourself in the center of that sphere.

And if you decide that for a water molecule at room temperature the sphere of neighbors with which there is a dynamic interaction has only a radius of 10 angstroms, whenever your unit cell is larger, is a cube, say, larger than 20 angstroms inside, you are studying, in practice, an infinite system because every molecule in that system sees all the independent neighbors that it needs to see, OK?

So even a simulation with a finite number of molecules can reproduce exactly the behavior of an infinite system provided a molecule interacts with other independent molecules up to a maximum range of dynamical correlation. And we'll actually define this quantity in the lecture that follows.

So again, length scales are another problem and a problem in which periodic boundary conditions tend to help a lot. But again, suppose that you want to study something like the strength of a metal.

Well, you have now a problem in which a lot of length scales becomes intertwined because in the strength of a metal, you might want to consider what is the dynamics of impurities that strengthen that metal. Like, if you have steel, you have carbon atoms in between the ions. And so you need to understand what is the atomic dynamics.

But then there will be dynamics of units on a larger length scale because in a metal there's going to be these location and slip and glide planes. And so there is coordinated motion of thousands of atoms that you want to consider.

And then, in reality, a real metal is never a crystal. But it's always a polycrystalline material. So you have grains. And inside that grain you have a consistent ordering of the atoms. But then grains are mismatched. And there are grain boundaries.

And so you have a lot of length scales to consider if you really want to predict what would be the mechanical response of a real system. OK, so this is the challenge, if you want, number 2.

The challenge number 3 is the one that we have described more in detail up to now that is determining how accurate can be your prediction of the interaction between nuclei. And you have seen everything that we could tell you about empirical potential quantum mechanical simulations.

And in principle, even if it's obviously very expensive to make very accurate predictions of the interaction between atoms, in principle it's the less conceptually challenging of the problems. I mean, we can just try all our electronic structure methods at will. And we can try to see how accurate we become.

And the last challenge that is also not really a conceptual challenge and it can be solved but is just a very expensive one to overcome is this one in which we treat the nuclei as classical particles. And that tends to be, in general, very good. But it's sort of an approximation that breaks down the more you go towards lighter and lighter nuclei because, remember, the de Broglie relation.

Your wavelength times your momentum is equal to the Planck constant. So the lighter you are, the smaller your momentum, the longer your wavelength, the more quantum you become.

So something like hydrogen that has a nucleus that is just one proton, so it's fairly light-- or it's the lightest of all possible nuclei-- tend to still have significant quantum properties even at relatively high temperatures. You see, the more you increase the temperature, the more the average momentum of your particles become and the more classical they become. But the lower the temperature, the more quantum they become.

And something like hydrogen is still significantly quantum at room temperature. And what does that mean? Well, it means that that particle, it's really a delocalized system that can actually tunnel through barriers. That's if you want one of the important differences between a particle that behaves according to classical mechanics and one that behaves according to quantum mechanics.

If you are behaving as a classical particle, you need to have enough kinetic energy to overcome a barrier. If you are a skier and you are in a valley and you want to go in the other valley, you need to have enough velocity to be able to overcome the mountain.

If you are a quantum skier, you can actually have some finite probability of tunneling through. And so there are tunneling problems that are important not only for hydrogen but sometimes even for much more massive particles or even for groups of particles, provided that the energy barrier that they need to overcome is small enough.

And I've shown you in the first quantum class an example in a perovskite crystal in which we have this cubic structure with an octahedral cage of oxygen that can configure itself in a ground state that is ferroelectric, in which the oxygen cage displaces with respect to the cubic symmetry and induces a permanent electrical dipole in your system. Well, there are certain materials that have this structure.

But since the octahedral cage can have roughly six possible choices on where to move offsite, if you decrease the temperature enough, the oxygen cage will start actually tunneling between all the six possibilities. Instead of finding itself at 0 temperature in a broken symmetry configuration in which the oxygen cage has chosen one of those six possibilities, in reality, because of quantum tunneling, it keeps tunneling between all the six possibilities.

And on average your system looks cubic. And on average you don't have really any ferroelectricity. But the polarizability of the system is going to be different from that one in which the octahedral cage was sitting squarely in the cube. So there is a dielectric response that is different.

So there are cases that are fairly exotic, I would say, for most applications, in which tunneling effects of the nuclei could become important. But there is one other important consequence that comes from quantum mechanics that tends to be important in a lot of cases. And that consequence has to do with the quantization of vibrational excitations.

That is, if you look at the motion of any system-- a molecule, a solid-- in classical terms, well, that molecule or that solid can have any amount of kinetic energy. In a hydrogen molecule, the atoms can have 0 kinetic energy if we treat it classically. Or they can move just a little bit. So they have a very minute amount of kinetic energy. And we can just increase the kinetic energy that those atoms can have at will.

In reality, because the system is a quantum system, the vibrational modes of a molecule or the vibrational modes of a solid are actually quantized. And you can't have arbitrary amounts of vibrational energy. Actually, in its quantum ground state, a molecule will not be at rest. But it will be in what is called the zero-point motion state. That is what is the quantum mechanical ground state.

You can think of the analogy with the stability of the atoms. In its ground state, the electron around the proton is not collapsing onto the proton because of the Coulombic repulsion. If the electron was a classical particle, it would collapse on the nucleus. And the hydrogen atom would be just a proton and an electron sitting in the same place.

But because the electron is quantum, it can't really collapse. And its lowest energy state is a state that is the 1s orbit for the hydrogen in which the electron is around the atom but doesn't collapse onto it. And the excitation of that electron to the next energy state is a quantized excitation.

You can't go to a level of 1s that has an energy for the hydrogen of minus 1 Rydberg to another state that has energy minus 0.99 or minus by 0.98 or any kind of amount. The next state in which the electron can live is a state that has an energy of minus 1/4 and then minus 1/9 and so on and so forth.

The same thing happens for the vibrational level of molecules, OK? A quantum molecule in its ground state is not in a state where the atoms do not move. But it's in a state where the atoms have a certain amount of quantum kinetic energy. And it's in a state called zero-point motion.

And if you want to increase the kinetic energy of this molecule, you can't do it continuously. There is a quantized cap. So you need to provide enough energy to heat up the system to the next state and then to the next state and then to the next state.

And that, if you think, affects deeply how you are going to heat up a system. That is, if you look at the specific heat of a solid, that is basically the quantity that tells you how much heat the system is going to be able to take, and you look at what happens at very low temperature, you discover that actually the system is almost unable to take any heat because all the vibrational excitations are much larger than the average amount of kinetic energy that you have if you are looking at what it takes to bring something from 1 Kelvin to 2 Kelvin.

OK, but this is still relevant if you have a system like water, even at 300 Kelvin. In water, say, at 300 Kelvin, what you have is you have these water molecules. Remember what you have-- your oxygen, hydrogen, and hydrogen.

And the modes of vibration of these molecules are stretching modes of the hydrogens. And then there are scissor-like modes in which the two bonds do this. And then these are intermolecular modes, just sort of relevant to one molecule. And then there are all the molecules interacting.

Well, at room temperature, all the internal modes of the molecule, the stretching and the scissor modes, have so much vibrational energy in their ground state that it's almost impossible to excite them to the next vibrational state with the average temperature that you have at 300 Kelvin. So even a system as fundamental as water has really a hybrid behavior in which the vibrational interaction between molecules can actually be populated arbitrarily by the temperature that is surrounding you.

But it's almost impossible to excite a molecule from a zero-point stretching motion to the next one. And so, say, the specific heat of water will be slightly exotic because really part of its vibrations are quantized.

And so there are a number of cases in which this classical nuclear approximation is not exact. And luckily they can be studied with specific techniques that we'll see in some of the later-application classes during this course.

But remember this-- if there is one thing that you want to remember about the behavior of nuclei-- that this vibrational excitation is quantized. And we actually see this quantization in physical properties like the specific heat.

OK, so let's discuss the first of the possible applications of molecular dynamics that I have mentioned. And that's using molecular dynamics to really calculate average properties of a system at finite temperatures.

And what I have written here is really the summary of statistical mechanics. So that is what, if you want, Boltzmann would tell you. You have a dynamical system that is described by this point moving around in phase space, OK? So with a certain possible set of positions and with a certain possible set of momenta.

And suppose that we want to calculate, say, what is an average property at a certain temperature. And we call that property A. It could be the lattice parameter of the system. Or it could be the diameter, say-- it's even simpler-- of your water bubble.

OK, you have this water bubble. And you want to see what is the diameter of this water bubble at a certain temperature. And this diameter is going to increase with temperature. And what you do, you just let this system evolve. And during its evolution, you keep measuring what is its diameter. And you let this evolve a lot. And you measure what is its width.

Well, the way Boltzmann would set this problem is say, we really need to integrate over all possible configurations of the system. That is, we need to go through all possible configurations of the system. That is, we need to integrate over the full phase space.

And for every possible point in phase space, we need to calculate the diameter of my water bubble. And then we need to take into account what is the probability of being in that configuration. And if you are in a system at a constant temperature, the probability of finding yourself at a given temperature is given by exponential of minus beta is just 1 over the Boltzmann constant times the temperature times the energy.

That means that configurations that have a high internal energy are very unlikely. And configurations that have lower energy are more likely, to the point that if the temperature is 0, that is this beta is diverging to infinity, only the lowest energy configuration possible is represented.

So this is the statistical mechanic points. To calculate an average quantity at a finite temperature, we need to sum over all possible configurations in space and momentum. For each possible r and p configuration, we'll have a certain value of the diameter of that bubble.

We weigh that with the statistical weight of that configuration. And we actually normalize this expectation value. And this is what we obtain.

So if, for a moment, you think at a phase space that is two dimensional and in this square we represent all the possible positions and velocities of our system, what Boltzmann tells us is that for every possible point, we will have a certain value of A . And the sum of all the possible values of A weighted with the probability of finding them gives us the expectation value.

And so what you really need to do is an integral over all space of this. And the cost of these integrals diverges as the number of dimensions diverge because, again, every time we increase the size of the system by one particle, we have six more dimensions over which to integrate.

And again, if you think that we do this just by numerical integrations by taking points on a grid and maybe we just discretize our system with 10 points along one dimension, every time you add a particle, you have 10 to the 6 points that multiply your system. So this quantity explodes.

The point of view of molecular dynamics is different. What we say is that, well, let's just let this system evolve. Let this system evolve in time. And let it evolve at a constant temperature.

And so now our point in phase space is going to go around, following a trajectory. And if we let it go around long enough, it will spend more time in the parts of phase space that are more favorite, the ones that are more likely, according to the Boltzmann factor.

And so if our evolution, really, is distributed in the trajectory according to this thermodynamic weight, thermodynamic factor, then our expectation value, the thermodynamic average for our observable A , is just trivial, obtained by integrating over the trajectories all the values of A . That is what I was saying before.

If you want to calculate the diameter of your bubble, what you do in the molecular dynamic sense, you just let the bubble evolve according to the equation of motion at a constant temperature. And you keep monitoring its diameter. And you will see that the diameter will have an average and that this is a much more efficient way than considering all possible configurations of the bubble, weighing them with how expensive that configuration is, and doing it for all the possible configurations.

The molecular dynamic algorithm automatically, if you want, keeps spanning the most likely configurations. And so it does, for you, the job of going in the place where it's most likely to be.

So in order to calculate a sample average, we can actually transform a problem into integration over a trajectory. And that is almost always doable under what is called the criterion of ergodicity. So the equivalence between the two formulations that I'd written before-- that is calculating a thermodynamic observable as an integral in phase space or as an integral over trajectories-- are equivalent, provided your trajectory is really able to go all over the place in phase space.

So if, again, this were phase space, you could have, in principle, the problem of a known ergodic system if there is something that prevents your trajectory to reach certain so-called forbidden segments in your phase space. And I can give you an example.

Suppose that you are studying something like silicon and you are in the crystalline state. And now you heat it up. And once you go above the melting temperature, your silicon will become liquid, OK?

Now, suppose that you cool it down again below, now, the melting temperature. OK, in principle, the liquid silicon should become again crystalline silicon because, at that temperature in which we have brought it back to, the thermodynamic stable state is really the crystalline solid. If you do that in practice with our molecular dynamic simulation, you see that you start from the crystal. You heat it up. It becomes a liquid with no problems.

And then you cool it back down. And instead of going back to a crystal, it goes back to a glassy, amorphous state. Silicon will find itself in a configuration in which locally all atoms are happy to be fourfold coordinated. But if we look at the long-range order, it's completely lost. The system is not crystalline anymore. It's glassy.

What's happening there? Well, it happens that it's just, for a liquid cooling down, so much easier to trap itself into a set of configurations in which atoms are really fairly happy as they are, because what is important for silicon atoms is to be fourfold coordinated.

But what is important to be in a crystal is that this fourfold coordination extends and becomes a long-range order. But for silicon, it's very inexpensive to mess up that long-range order, like for many other systems.

And in principle, same thing happens for polymers. You lower it down. And they sort of get trapped in a glassy configuration. A lot of polymers might want to stay actually in a crystalline state. But really that is impossible to reach.

And so if you want to study, say, what are the properties of your silicon and you have obtained or you are trying to obtain your silicon structure from a liquid that you cooled down, you'll actually-- in your molecular dynamics, you'll keep orbiting around in this glassy mass state and will never reach that region of phase space that is actually very, very small because it needs to have all atoms long range in very specific positions.

And so you are really spanning a system in a nonergodic way. That is, you spend a lot of time in a region of phase space that are not really representative but out of which is very difficult to escape.

OK, so there are problems in, again, brute force molecular dynamic approaches that you have to be careful about. OK, so let's see how we actually go and try to do this on our computer experiment. And I'll start with the simplest case that is what is called the microcanonical case.

That is, we perform a simulation where the number of particles is constant, the volume is constant, and the total energy of the system is constant. This is what is called a thermodynamic ensemble. Again, these are all extensive variables. And we are fixing them.

Sometimes it's appropriate, instead, to do a Legendre transformation and go into a different thermodynamic ensemble. Say, you could want not to be at a constant volume. But you might want to study a system at a constant pressure where the volume changes.

And so we need to describe how we'll actually be able to do that. That is, how will we be able to move to a system in which different intensive thermodynamic variables are constant?

But if you take a system of particles in periodic boundary conditions and you have 10 water molecules in a certain unit cell and you evolve them with the Newton equation of motion, you can obviously know that the number of molecules and the volume of the unit cell are not going to change. And if you integrate your equation of motion appropriately, the total energy of the system, that is the kinetic energy plus the potential energy, is not going to change.

And so as we evolve this point in phase space, its trajectory is going to span what we call the microcanonical ensemble. And you let this point evolve long enough. And all of a sudden, you will have a meaningful trajectory, that is a meaningful assembly collection of microstates of representative configurations in space and in velocities of your system, out of which you can calculate averages of any of the observables that you want to calculate and that you want to characterize.

So what do you do in these molecular dynamic simulations? Well, there are really four important steps that you need to take into account. And we keep thinking of water, liquid water, as an example. That is, you need to start from something. That is, you need to put your water molecule in a certain position. And you need to give them initial velocities.

And you are almost never going to be good enough in choosing a state of the system that is somehow meaningful. You'll always, when you start building something complex, put molecules in places that they don't want to stay. And one needs to be very careful, then, in all of this.

But suppose that we have, in any case, an initial configuration of position and velocity. Then what we need to do is, over and over again, integrate the equation of motion, that is, from the current position and the current velocities, predict the new position after a certain time.

And once you have the new position at a certain time, you can recalculate what are the forces acting on the particle in the new position. And with these forces, you can, again, move the particles by a certain amount and so on and so forth.

And that is where the molecular dynamics algorithm comes into play. We need to integrate our equation of motion. And with this, our system will start to evolve.

We have put that point in phase space in motion. And it'll start moving around. And if we are integrating correctly, the total energy of the system is going to be conserved.

Now what we need to do is we need to make sure that our characterization, our thermodynamic averages, do not depend on our initial condition. Again, we have started from a configuration that we have built by hand. And because of that, it's most likely a very unfavorable configuration of the system.

OK, so we need to let the system evolve and somehow start figuring out really in which condition it would want to stay. Think of the previous example of silicon that is glassifying. Well, we had liquid silicon. We have to cool it down.

In a way, what's happening is that we are never equilibrating it because the system goes into a glass phase. And so it never loses memory of its initial condition. It's always sort of trapped in this pseudo liquid, now glass, state. While in reality, at a temperature below the freezing or the melting point, equilibrium would be given by a crystalline state with the ions vibrating.

So there are cases in which equilibration is so difficult that we actually never reach it. Silicon never goes to a crystal. That's why we need to be careful. But in most cases awfully we are lucky. And we lose memory of initial condition.

Say, the opposite is actually very easy to achieve. If we start from the solid and we warm it up, it becomes a liquid. And if we wait long enough, that liquid will have lost all memory of having come from a solid.

OK, so it's actually very easy in one direction and very difficult in the other direction. But once this equilibration time has passed, well, at that point we have truly everything in place. Our point in phase space is moving around. And it's really going into the regions of high probability and distributed, according to Boltzmann distribution.

And so at that point, we can start an average among any of the physical observables that we want to calculate. And so there are these four phases. And I'll briefly mention how we actually choose them.

Initialization. Well, again, we need position and velocities. You need to be careful, in particular, to avoid overlap between your molecules.

If you put two atoms very close, they are going to repel each other with an enormous strength. And if you put them close enough, they are going to shoot in opposite directions very quickly and completely kick off and mess up everything on the side.

So what is difficult is really, always in configuration space, finding the right structures. This is the difficult part of choosing initial condition. Choosing velocities is much easier. One possibility is just giving zero velocity to the system or a very small velocity and then slowly adding things.

And what we usually see is that configuration in velocity space is much faster. And this is the reason. Think of that trajectory for a moment of one particle. And think of what happens when you have a collision.

Suppose that you have particles, for a moment, that are hard spheres. In a collision, the velocity of the particle all of a sudden changes sign. So in the velocity part of the phase space, your point makes a sudden jump, OK, because the velocity that was plus 10 has become all of a sudden minus 10.

So in velocity space, for the limit case of a system of hard spheres, your representative point hops everywhere continuously. In reality, I mean, you have interaction at all times. And you don't really hop. But you can sort of move very quickly from one region to the other region.

There is not anything similar to the glass entanglement in configuration space. So it's very easy to thermalize velocity. It's very difficult to thermalize and equilibrate position. And that's where, all, you need to be careful.

And if you really want to be accurate, what you can also do instead of giving zero velocity to your system, you can use your concept of statistical mechanics. That is, a system of classical particles will have a distribution of velocities that depends on the temperature at which it is. And this is 0 in Celsius. So it's 273 Kelvin.

And the distribution of velocity in a classical system is actually given by the Maxwell-Boltzmann distribution written above that has basically this shape. So it means that if we are, say, at 0 Celsius, 273 Kelvin, we'll have a lot of molecules with a certain velocity. And then, in principle, we can have molecules with higher and higher velocity. But the probability of finding them decreases.

And as we increase the temperature, the probability distribution of the velocities for the molecules changes. And so as usual, you can define both a median or an average in your system. That is, you can look at what is the most likely velocity, that is the top of this curve, at 0 Celsius. Or you could also define what is the average velocity. And since that this curve is asymmetric, it's going to be really shifted a little bit.

And these are actually the numerical expressions for these two velocities, the most likely and the average velocity. And so I've given you the example of an oxygen molecule at room temperature.

We are talking about 1,000 meters per second. So these velocities, if you want, are similar to the speed of sound in your material because, at the end, sound propagates at a certain speed because that's the velocity of the collision from one molecule to the next that are really carrying sound away.

So if you want to be sophisticated, you give to your system a set of velocities that could be, say, all identical to the most likely velocity. But again, equilibration and thermalization in velocity space is very easy. So once you have this, that is, once you have a set of positions and the set of velocities, you need to use Newton's equation of motion of saying, of predicting, where your system is going to go.

And so that's when the molecular dynamic algorithm comes into play. So what you need to do is, as we say, you need to use an integrator, the algorithm that tells you what your next position and what your next velocities are going to be.

And there are a number of integrators. There are subtleties. But often you find the name of Verlet algorithm. This is one class of algorithms. And the other common class is the Gear predictor-corrector, just to give you the terminology. And we look, in a moment, at the Verlet algorithm.

They're all fairly simple. But what is really important is that this algorithm should be robust. That is, they should give you a trajectory that is very close to the trajectory that you would get from the perfect analytical integration of the equation of motion.

And in particular, what is the most important thing is that they should be accurate in, say, conserving the constant of motion. As we have seen from the Newton's equation of motion, the sum of the kinetic energy and the potential energy should be a constant.

If your integration, if your numerical evaluation, of that equation is poor, you will see that the constant of motion is not any more constant. So again, this is the sanity check on the accuracy of your integration.

And there are more subtle elements of the integration algorithm that can become important. But we won't go into that. And then as I said, straightforward integration of the Newton's equation of motion gives you the microcanonical ensemble in which the total energy's conserved.

But sometimes you want to study systems in which the temperature is fixed. Or you could want to study systems in which the pressure is fixed. And so you need to augment your molecular dynamic integrator with some feature that allows you to control those intensive variables.

So if you want, say, to control the temperature of the system, what we say usually is that we couple the system to a thermostat. That is, we make sure that the temperature in your evolution doesn't change. And you'll see examples in the next class of this.

And in the case, in particular, of something, say, like the temperature, you have a whole set of possible thermostats. That is, you could use stochastic approaches. That is, you have your evolution of your system. And your system doesn't conserve, if it's microcanonical, the temperature. Sometimes it goes towards higher temperature. Sometimes it go towards lower temperature.

And what you could do is either every now and then absorb some energy or give back some energy in your system-- and that would be a Langevin dynamics-- in order to control the average kinetic energy of your system. Or you could somehow forbid, in the microcanonical evolution, your system to choose a state in which the temperature changes.

So if your new positions are such that your velocity gives an increase in temperature, you just shorten how much your atoms move in order to make the velocity smaller and in order to make the kinetic energy conserved. Or there are more complex dynamical ways of imposing temperature that goes under the name of extended Hamiltonian, or extended system, of which we'll see an example that is [INAUDIBLE].

But let me give you really what is the simplest and still one of the most used numerical integrators. And it tends to be very robust and very accurate. So it's still used in a lot of molecular dynamic simulations. And there is often no reason to do anything better than this.

So in a way, we are in a much luckier situation than in quantum mechanics. This is good enough in almost all cases. And this is how we look at the problem. That is, what we need to do always is we need to calculate the position of each particle at a time $t + \Delta t$ once we know the position, the velocity, and the forces on that particle at time t .

So suppose I'm at a certain instant in time that I call t . I know everything about my system-- position, velocity, and forces. What I really need to do is predict what are the new positions. And when I'm in the new positions, I want to calculate the forces again and evolve the system.

And now, what do we do? Well, let's do a simple Taylor expansion. That is, this is r of t is a function of the variable t . So we can use a Taylor expansion. And r of t plus Δt is going to be r of t plus the derivative of r with respect to t times Δt .

And that's the velocity. Plus $1/2$ of the second derivative of the position with respect to time. And that's the acceleration. Times Δt squared, and so on for the third derivatives and so on for the fourth, fifth, ad infinitum.

This is just the Taylor expansion as a function, OK? And so it gives us the position of t plus Δt , knowing all the derivatives-- first, second, third, fourth, n th-- at times t . And once we have the derivatives, we can not only predict what the position would be at t plus Δt , but also what the position would be at t minus Δt . And it's written here.

And if you think, the terms with an odd power change sign in Δt . And the terms with an even power in Δt do not change sign. So we have just done a Taylor expansion. And now what we do is we actually sum these two quantities together, OK?

And this is what we obtain, OK? We sum the two quantities together. All the odd terms disappear. This sums to 0. This sums to 0. The fifth order, seventh order, and so on sum to 0.

So what we have is that we have r plus Δt plus r minus Δt is equal to 2 times r plus α -- sorry-- acceleration times Δt squared. It's just written here. And then rearrange the terms.

But you see, this is the Verlet algorithm. Now what we have when we look at this expression, we have an expression that tells us what is the position at my new time step, knowing the position at my current instant, knowing the position at the previous instant.

We are keeping configurations from our trajectories. So we know where we were before. We know where we are now. And we know where we are going to be just by adding to this the term acceleration at the present time times Δt squared.

We make an error because we are not introducing terms that are higher order in Δt squared. But you see, the term in Δt cubed has canceled out. So the error that we make is just of the order Δt fourth. That is, that means that the smaller we make the Δt , the smaller we make the time step, the more accurate our integration is.

And so this is, if you want the most critical parameter in your molecular dynamic simulation, you need to make sure that the steps that you take are small enough, considering what is really the variation of your force as a function of time because remember that now that we use Newton's equation of motion, the acceleration at a time t is just given by the force divided by the mass. And because the force is just the gradient of the potential, it's just given by this.

So at a certain instant in time t , you have the acceleration acting in your particles just by taking the force and dividing by the mass. So at every instant, you calculate force. And the force lets you evolve the system. And it keeps giving you the new positions.

And so you know what is the trajectory and configuration space. And you can very simply find out what is your velocity at every instant t , if you need to calculate what is the velocity just by doing a finite difference calculation of the tangent basically. The velocity is the tangent of the trajectory with respect to t .

OK, this concludes class. What we'll see in the next lecture is how we actually take this Verlet integration and how we use it to do the characterization of the thermodynamic properties of a real system.

So today is Thursday. Next week is spring break. So we'll see each other again on Tuesday, 29th. And I will be at a conference. And all the TAs will be at a conference next week.

So it will be slightly difficult communicating via email. So try to ask a lot of questions between now and Sunday for your problem set that is due the Tuesday after spring break. And that will be problem set 2. Otherwise, enjoy spring break.