

**NICOLA
MARZARI:**

OK. Welcome, everyone, lecture 12, second part of molecular dynamics, also after the spring break. So I decided to spend two or three slides in reminding you what was going on from last time. And I've actually been intrigued.

I'm starting with Heraclitus because this is a beautiful citation. So this is, you know, fifth century BC. But there is the intuition that a macroscopic property, like the river, actually comes from a microscopic property.

That is the water molecules. Though he wouldn't know what a water molecule is, but comes from what the water molecules are doing. So Heraclitus was saying, you cannot step twice in the same river. Because the molecular conformation is always different, but the river is always the same.

Anyhow, reminder of what we have seen last time-- first of all, of course, the standard thing is that we have looked at Newton's equation of motion. Really, molecular dynamic simulations involve the integration of these couple set of second-order differential equation. So if we have n particles, there are going to be n trajectory defined by, for each particle, how the vector position evolves with time.

And the force acting on each particle can be obtained from the derivative of the potential field. And all the particles interact with each other. So this is still a many-bodied problem, although a much easier one to solve than the electronic problem.

And we sort of remind you how we integrated this equation of motion. But sort of one of the fundamental concepts that we have seen is how we can use actually the integration of the molecular dynamics equation of motion to obtain thermodynamical quantities. And you see what I'm writing here on the left is really the Boltzmann idea, the sort of basic idea of statistical mechanics.

That is, if you have a certain variable that I'm generically calling here A , what we know is that the expectation value for that thermodynamical variable is given by the integration over all the phase space, integration over all the position, integration over all the momenta of the value of that variable A for every position and every momentum weighed with the appropriate probability distribution.

So this would be, say, for what we call a canonical distribution, a distribution in which each microstate, its single position and momentum, is weighed with the exponential of the energy with a coefficient β that is just 1 over the Boltzmann constant times the temperature. And this fundamental sort of statistical mechanics approach can actually be translated into a molecular dynamics approach.

That is we can actually obtain the same expectation value for the operator A without really doing these integrals. These integrals are obviously impossible to perform explicitly as soon as you have many particles because the dimension of your phase space explodes. And we have seen it over and over again that, as the dimension increases, it becomes increasingly difficult or computationally impossible to perform these integrals.

But what we can actually do is evolve our system. But, remember, our system is just a point in phase space, the collection of all position and momenta. We evolve it with the molecular dynamics equation of motion. And in particular, say for this example, we need to evolve it in what we call the canonical ensemble.

That is we need to evolve under the condition that the temperature of the system is constant. But in this trajectory, in this evolution, we can average for each instant of this trajectory what is the instantaneous value of your observable A. And the average over the trajectory gives us an expectation value.

And these two quantities are equivalent. They are obtained in two different ways. Here, as an average over the trajectory, that is much, much simpler to do computationally. And here is an average of all phase space.

And there is a subtlety about this equivalence that is called the hypothesis of ergodicity. That is, in order for this equivalence to be valid, we really need to think that our system spans the phase space with the appropriate probability distribution without being confined in corner of phase space. And I'll give you an example later on of what means being confined.

So in a way, we sort of work under the hypothesis that really this trajectory is spanning an appropriate sample of trajectories. And so when you do a molecular dynamic simulation-- again, reminder of last class-- you really perform an ideal computational experiment in which you start deciding what are the position and the velocity of all your particles. And then this is sort of the computationally important step. You integrate the equation of motion.

And this is the sort of step in which you need to pay particular care and attention on the algorithm you use to integrate your equation of motion and, also, on the accuracy that is needed to integrate this equation of motion. And once you have put all your particles in motion and you are sort of integrating them and evolving them according to the appropriate equations, well, you just need to make sure that you wait long enough. That is that you equilibrate your system, so that you basically lose memory of your initial condition.

You see, your system starts from a choice that has been made by you and by hand. And in data, it's never representative of what could really be a microscopic configuration of your system. You're putting molecules certain position, and you're giving them velocities. But that's just your assumption.

And so you need to lose memory of this initial condition. And that's accomplished by evolving the system long enough. And how long is the time needed to reach equilibrium is something that depends on your system, depends on the temperature at which you are performing a simulation. But it's something that you can check in the course of your simulation.

And again, I'll show you an example on how you check for appropriate equilibration. And once you have reached equilibrium, well, then it becomes straightforward to just calculate the results of your experiments. That is calculate, say, what are the thermodynamical averages. And so you start accumulating what are your quantities of interest.

And again, sort of reminder of last class, the initialization step involves choosing the initial position and the initial velocities. And we need position and velocity because really we are dealing with second-order differential equations. It's more critical to choose appropriately the initial position and more difficult.

If you are studying a system that is crystalline, obviously it's very simple. You start just with sort of a periodic arrangement of atoms. But say if you are starting something more complex, like, say, liquid water, it becomes more critical. Because your liquid, as a structure, that is not immediately evident to you when you try to put molecules around.

Anyhow, there are sort of obvious things to avoid. That is you want to avoid that atoms overlap too closely. Because if you put two atoms too close together, they will have an initial repulsive force that is very strong. And they'll just zoom away very quickly from each other.

It's much more critical to equilibrate with respect to position than with respect to velocities. Velocities, or if you want the phase space part, the momentum part of your phase space, is much easier to integrate. Basically, because in any collision between atoms or between molecules, the velocity, if you want, get reversed.

And so in the momentum, in the velocity part of phase space, your representative points can really jump around. Say, if you had an ideal gas of hard spheres that only collide when they touch each other-- and in an elastic collision, they just reverse their velocity-- well, in that case, your representative point is really jumping around discontinuously.

So in velocity space, you really move around very easily. And you span all the phase space. It's really in the position part of the phase space that you need to be careful and where your system can become entangled.

Once you have set up your initial position and your initial velocity, you have everything you need to integrate your equation of motion. And to integrate equation of motion, there are a variety of algorithms. And I have mentioned, and you'll see again, the standard or simple Verlet algorithm.

That is actually a very robust and very accurate algorithm. So it's what you'll see in the next slide, but there are sort of different varieties of Verlet algorithm. Sometimes you see them called leapfrog Verlet or velocity Verlet depending on the specific integration scheme.

And then there is another class that I'll also mention briefly just for reference that are called the Gear predictor-corrector integration algorithm. So this, more or less, relies on all the zoology of integration algorithm.

And obviously, one of the important qualities of your integration algorithm is that it must be robust. That is it shouldn't be too sensitive on the choice of discretization. Remember, we have sort of differential equation that are basically perfect analytical equation.

But in order to integrate them on a computer, we need to discretize them. And again, sort of one of the key features of your simulation is figuring out what is the discretization interval or what we call the time step. That is how much time you can spend before you calculate another point in your trajectory.

And there are sort of a lot of subtleties that we won't go into on what really makes a good integrator for the equation of motion. In particular, one of the properties that we would want is that it's time reversible. And that's what the simple Verlet is.

That is, if you invert all your velocities, your trajectory is recovered in the opposite direction. And not all algorithms actually are time reversible. And then there is something a little bit more subtle. That is a sort of criterion of robustness for an integrator is that it preserves the volume of your representative system in phase space, but we really want to go into these details.

There are sort of, you know, several classes of molecular dynamic simulations depending on what is the thermodynamical ensemble that you want to represent. That is, if you straightforwardly integrate Newton equation of motion we have seen in the last class, you actually conserve the total energy of the system. That is you conserve the sum of the kinetic energy plus the sum of the potential energy.

And as you know from statistical mechanics, that is what is called the microcanonical ensemble. Sometimes it might be more worthwhile to perform your simulation at a constant temperature. And so that is what you would call the canonical ensemble in statistical mechanics.

And so you need to find a way to impose a constant temperature in your system. Your system, in principle, would conserve the energy. And so it would have a fluctuating temperature.

And again, from statistical mechanics, the larger your system is, that is the closer you get to the thermodynamically limit, to the limit of an infinitely large system. It really doesn't matter because an infinitely large system that is simulated with, say, a microcanonical algorithm will also have 0 fluctuations on the average temperature. And so you would obtain the same results in a microcanonical or in a canonical simulation.

But as always, we deal with finite systems maybe replicated with periodic boundary condition. And so it might be more convenient to choose a canonical simulation, so that we can fix the temperature of our system. Suppose that, again, you want to study water. You want to make sure that you are studying your system at a temperature that is consistent with liquid water.

Suppose that you were studying water with a microcanonical simulation. Again, what you would do is choose your initial positions, choose your initial velocities, and let your system evolve. But it's only after you have let your system evolve for a long time, that is after your system has equilibrated, that you can actually measure, start averaging, some of the macroscopic properties.

And so in the microcanonical ensemble, you could start, after a while, to average as a macroscopic property the temperature. The temperature is nothing else than the average kinetic energy of the ions. Well, suppose that you are studying water. You have set up your initial position, your initial velocity. You let this evolve.

And after a while, you discover that the average temperature is 150 kelvin. And your system is frozen. OK. Well, what you have to do is then restart your simulation. It's not going anywhere if you want to study liquid water.

And how could you change this? Well, you could try using the same initial configuration that you had used before, but much larger velocities, hoping that once the system evolves deterministically from that initial point, your average temperature is going to be a temperature in which your water is liquid. But maybe you do it again, and your velocity is too large. And now, the temperature of your system is 500 kelvin. OK.

So you see the problem, in this case, in starting with a microcanonical approach. So it would be much more convenient to have an integration scheme that actually allows you to fix the temperature instead of allowing you to conserve the energy. And there are several ways in which something like this can be done.

And again, you'll see examples. But generally speaking, we say that we couple our dynamical system with a thermostat. That is nothing less than, if you want, a computing device that allows us to control, in this case, the temperature of our simulation.

And if instead of the temperature we wanted to control the pressure, we would couple our system with a [INAUDIBLE] and so on and so forth. But if we focus for a moment on the thermostat that's the simplest, well, there are sort of several ways in which we can control the temperature of our system. And I've mentioned here sort of three classes of devices that we use to couple our system with a thermostat.

And one approach could be the Langevin dynamics or the stochastic approach. That is you try to sort of simulate a constant temperature by letting your system evolve with the microcanonical equation of motion. But every now and then, you give little kicks to random atoms in which those atoms either accelerate-- so they start picking up more velocity. And so the temperature of your system rise. And so you drive your system to your required temperature.

Or instead of sort of kicking and accelerating them, you slow them down sort of with a kick in the direction opposite to its velocity. And so you cool the system down. And this is really sort of analogous to the thermal bath if you want. You have all your molecules moving around, and collisions thermalize them.

So this would be a stochastic or Langevin approach. Another approach would be a constrained approach. Often, the simplest one is the velocity rescaling. Suppose that, again, you want to keep a constant temperature.

Well, constant temperature means that you have fixed overall average kinetic energy for your ions. And so when integrating the microcanonical equation of motion, at every time step you have the velocity on every atom. And if the average kinetic energy is below your target temperature or above your target temperature, the only thing you need to do is change your velocity in order to have the constraint that, at every time step, the total kinetic energy of the ion is fixed.

So if you want, a constrained thermostat involves keeping the total kinetic energy of your system fixed. And in order to do this, you need to have actually an integrating algorithm that uses the velocities to integrate the equation of motions. So something like the leap frog Verlet that we want, see, or the velocity Verlet that we want, also see, would work.

But the simple Verlet that I'll show you again in a moment wouldn't actually be useful. Because in the simple Verlet, you never control or never use the velocity of the system to integrate the equation of motion. So you are not able actually to control the velocity of your system. And so you can't thermostat your dynamical ensemble using the simple Verlet.

And sort of the last approach that is probably the more generic and the more powerful is actually adding to your sort of set of equation of motion an additional dynamical variable. So up to now, the dynamical variables are all the position and all the velocity of your particles in the system. Well, you could add an additional dynamical variable that I tend to sort of compare it to some generic wind going through the system.

And that sort of additional dynamical variable is something that either slows down your particle or accelerates your particles in order to fluctuate appropriately around your target temperature. And this is the most elegant sort of solution to the problem of the thermalization, and it's also the most powerful. And we'll see examples both in this class and especially the next class of applying this extended approach.

And often, it goes under the name of Nosé that was the first one to develop the equation of motion for a canonical simulation at constant temperature, or actually Nosé-Hoover. That was the other person that made this equation a little bit more tractable.

And again, a reminder of the simple Verlet, the only algorithm that we have seen, that is sort of summarized in here. And what we are doing is we are looking at this trajectory, this vector that evolves as a function of time. And we are just doing a Taylor expansion.

That is, if we know the position at time t , the position at times t plus Δt , the Taylor expansion, involves the derivative of the position with respect to time. The second derivative of the position with respect to time, that is nothing than the acceleration, the third derivative and so on and so forth. And I've only expanded it up to the third order.

But the very nice thing is that we can expand it both forward in time and backward in time around the instantaneous position. And so using, you see, the instantaneous velocity, the instantaneous acceleration, and the instantaneous third derivative, we are able to correlate the present position with the next position in time and the previous position in time where Δt is this fundamental discretization unit that we call the time step.

And that is the unit that we choose to integrate our equation of motion accurately enough. And just by taking these two expression and summing them, we see that all the terms cancel out. And just by sort of moving left and right the appropriate term, we have just, by summing these two Taylor expansion, there's a fundamental expression that is really the core of the simple Verlet algorithm.

And you see, this is the recipe that we use, actually, to integrate our equation of motion in time. That is what we are saying is that the position at times t plus Δt is related to the position of time t and the position at times t minus Δt . And so by knowing this position and the previous position and by knowing what is the acceleration at the present time-- and the acceleration is nothing else than the force divided by the mass.

Here is where Newton's equation of motion comes into play. We are able to predict what is going to be the next position with an accuracy that is of the order of the fourth power of your time step. So if you want, if you are going to have your time step, the accuracy that you obtain is going to be 16 times higher.

So one of the fundamental sort of checks that you want to do in your molecular dynamic simulation is that your time step is small enough. But obviously, it doesn't need to be too small. And again, we'll see an example on how you can check for this.

And you see the simple Verlet algorithm never uses the velocity of your system. It's able to calculate the new position just using the present acceleration of the present force. And so it's not an algorithm that allows you to control the velocity of your particles.

So the velocity of your particle is something that is given to you and you can, of course, calculate. Suppose that you want to calculate what is the kinetic energy of your system. Suppose that you want to find out what is the average temperature.

Well, you need to calculate the velocity. And again, you can use a simple finite difference formula. And you see, what we use is actually the position of t plus Δt that's written here and the position at t minus Δt that is written here.

And these are symmetric derivative. And you could actually do the homework of inserting these quantities and finding out what is the order of error in a power of Δt for your velocity and why we use, actually, this symmetric formula, say, compared to a left or a right derivative. You could also calculate the velocity as $\frac{r(t + \Delta t) - r(t - \Delta t)}{2\Delta t}$.

And it's easy to do it here. If you want to calculate what is the velocity at times t , you obviously need to know what is the derivative here. This is the definition of velocity.

And you can see that the most accurate definition of a derivative is obtained by taking the value at $t + \Delta t$ and $t - \Delta t$ and measuring this slope rather than, say, measuring this slope or measuring this slope. So this slope here is given by the symmetric derivative that I have written down here and is the sort of most accurate representation of the true slope that is the green position here.

And you can also do this not only sort of graphically, but analytically, just writing out the explicit expressions. I won't go into the details of other integration algorithm. But just to mention another very common one is what is called the Gear predictor-corrector algorithm.

This is actually not time invariant. And it tends to be probably more accurate than the Verlet algorithm if you have an analytical expression for your forces, so if you can actually derive from a potential field your force and then obtain your acceleration. It tends to be much worse than the Verlet algorithm if you have noise in your data.

And that's actually the case when you perform an ab initio molecular dynamic simulation because the forces that you obtain in an ab initio simulation are exact only, if at each time step, you have reached the perfect self-consistency.

And so ab initio molecular dynamics has some inherent amount of noise that calls for the use of a more robust integration algorithm, like this simple Verlet. But if you are doing classical molecular dynamics, this is often superior. And really it's sort of an extension of the Verlet algorithm.

That is, instead of using just the first and the second derivatives of your trajectories, it uses the third, the fourth, and the fifth derivative, and so on and so forth. And then sort of it estimates the error and corrects, once you have taken your time step, for the difference between your predicted, say, acceleration or third derivative or fourth derivative that comes from your Taylor expansion and actually, say, the acceleration that you calculate. That is the force that you calculate once you move in that position.

And again, using higher order expressions-- that is using third and fourth derivatives-- can give you more accuracy and allow you to use longer time step to integrate appropriately your trajectory if you have no noise in your forces. But if you have noise in your forces, it can actually be more catastrophic in the errors that it obtains. OK. This was a reflection of some of the things that we have seen.

One of the sort of very simple concept I wanted to sort of remind you that takes place whenever we try to do numerical integration of a differential equation is the concept of Lyapunov instability. That is sort of what you have heard over and over again called the butterfly effect in weather forecast. That is, if a butterfly beats his wings in northern Africa, we have a tornado in the Caribbean.

And the reason for that is saying that, in any of this differential equation, when you try to integrate them, an infinitesimal error at the beginning of your integration will exponentially increase to the point that what comes at the end of your simulation can be radically different for infinitesimal differences in your initial condition.

That is to say suppose that you have an exact analytical trajectory. That is you are able to integrate your equation of motion starting from a certain initial position and initial velocity here. So in the limit of an infinitesimal time step, in principle you are going to do a perfect integration and recover this trajectory.

In practice, it means that, for larger and larger steps, what you will have is that you'll be able to sort of follow this trajectory for a certain time. But after a while, you'll start diverging. And the smaller your time step is, the more accurately you are going to reproduce this trajectory for a while.

But at the end, the exponential explosion will always take place. That is you start having an exponentially increasing difference between the exact trajectory and the trajectory that you are calculating. Luckily, this doesn't really matter.

This is very important. And this is sort of true in molecular dynamic simulation of our small sample periodically repeated in, which we want to calculate thermal averages. And it also applies to the case of the butterfly beating its wings and creating a tornado.

What really happens is that all these systems on the long time scales are basically inherently chaotic. So all these trajectories diverge from the ideal trajectory, but it doesn't really matter. When we do a molecular dynamic simulation, we are not trying to reproduce an exact trajectory starting from the atomic timescales and propagating for picoseconds, nanoseconds, or microseconds.

Because what we want to find out are either average properties of the system-- well, it doesn't really matter that the trajectory diverges. Because even this trajectory that diverge still conserve the fundamental macroscopic variables. If you are doing a simulation at constant temperature, what it really matters is that all these trajectories correspond to evolutions at constant temperature. And then you are sampling appropriately your phase space.

And even if you want to do molecular dynamics of an event instead of using it to obtain an ensemble average, even if you want to, say, find out how an atom diffuses across a barrier, what is really your macroscopic observable is the average over tens, thousands, millions of these events, all sort of slightly different from each other, but all corresponding to the overall constraint of having the same macroscopic observable.

So it doesn't really matter that we replicate exactly the trajectory, neither for an ensemble average of thermodynamical quantities, nor for the case in which we simulate a specific event. Because in any case, what we need to do is repeat the simulation over and over again to have appropriate statistics and to have an average. Of course, the only case where having the exact trajectory matters is that if you are sending a satellite in outer space because you are not sending 1 million satellites and seeing what they do on average.

You want that satellite to do something very precise. And so you need to integrate things very accurately. And you need to keep correcting your integration as time goes by. And we can actually sort of formalize this sort of concept of Lyapunov instabilities in terms, basically, of an exponential divergence in some metric of our actual trajectory from the ideal analytic trajectory of our system.

OK. So now, everything is in place. We are doing our molecular dynamic simulation. First of all, how do we check that something is working, that what we are doing makes sense?

Well, the first thing is just making sure that what you are doing makes sense computationally. It might still not make sense physically, but your sort of order [INAUDIBLE] is making sure that you are integrating your equation of motion appropriately. Because, remember, you need to choose a time step. You need to choose a discretization interval.

And the size of that interval depends really on what is the typical atomic motion. If you are studying a system that is very cool and is made of very heavy atoms, you can use very large time steps. If your system is very hot and made of very light atoms that move much faster, you need to use very small time steps.

But on average, there is really a small range of time steps that is really acceptable. And they are of the order of the femtoseconds, 10^{-15} seconds. So for something as light as a hydrogen atom bound to an oxygen, say, in water molecule, you might need to use a time step as small as $1/2$ femtosecond.

For something in which maybe you are studying, say, liquid caesium, a free electron metal that is very heavy, you might have a time step that can be appropriate and as large as 20 femtoseconds. But this is the range, basically, of the order of femtoseconds. And the very important thing is that you can always check that your time step is accurate enough.

Because if you are doing a microcanonical simulation, what you need to have is that the total energy-- that is the sum of the potential energy and the kinetic energy-- is conserved or, if you want, that the kinetic energy of the ions and the potential energy of the ions mirror each other exactly. And again, in the limit of infinitesimal time step, this total energy line should be perfectly flat.

As you increase your time step, well, this line will start to have some microscopic oscillation. So it will start fluctuate and, depending on your algorithm, could also have a systematic drift. A systematic drift, generally speaking, is very bad because it means that your system is really cooling down.

That is a systematic drift means that your total energy is going to get closer and closer to the potential energy. So the kinetic energy of your ions is decreasing and decreasing. And so systematic drift means that, after a while, your system will slow down to a halt.

And so you need to check for that. And then even if there is no systematic drift, a simple algorithm like the Verlet tend to be very good at that, at not having any systematic drift. You need to make sure that your fluctuation that will always exist as long as you are using a finite time step are actually small enough. So you want the fluctuation, your total energy, to be much, much smaller than the fluctuation in your potential energy, say two orders of magnitude, three orders of magnitude smaller. So that locally your system is varying its potential energy, but its total energy is actually unchanged.

And again, thanks to the power of the Verlet algorithm that sort of scales as the fourth power of your time step, you can do very simple tests. You half your time step, and you should see your total energy curve becomes much flatter. And so you can basically choose sort of what is the optimal time step.

And you can even use a sort of real space picture in which you can look at the trajectory. After all, what you have are atoms moving around and oscillating. And so what you want to do in a real trajectory of an atom moving around is you want to have enough points to describe correctly, say, a period of an atom.

And the rule a thumb is that you want between 20 to 100 points in describing a full period of an atom. And that, again, sort of brings us back to the order of magnitude of femtoseconds for the time steps of your systems involved. So once you have set up your computational part correctly, you still need to check out that your simulation makes sense.

And in particular, the sort of trickiest part is making sure that you reach thermodynamic equilibrium, so that you lose memory of your initial conditions. And again, that's actually something in which the Lyapunov instabilities, if you want, help. No matter how bad our choice is, if we let the system evolve long enough, our trajectory will sort of diverge from our initial exact trajectory that was somehow poorly chosen.

And it will become more and more representative of the thermodynamical ensemble that we are simulating. But it's still very tricky to understand if we are reaching equilibrium. I've shown you here an example in the case in which, say, we are looking at a solid, something like a slab of metal in which we have a, say, bulk atom in the metal.

And then as we move towards the surface, we have sort of second layer atom and first layer atom. And what we are averaging here, what is the sort of thermodynamical observable we are interested, is the mean square displacement. That is how much the atom fluctuates around the initial condition. And again, you'll see more and more example in what comes later of a mean square displacement. But you see, somehow it's written here in a reverse scale.

But if we start at time t equal to 0 and we measure what is the mean square displacement, well, at the beginning, the atoms will keep moving. If you think, this peak here, this sort of $1/3$ to $1/4$ of a picosecond, is a measure of what is the period of oscillation of an atom. Because at the beginning, the atom starts moving.

And from the point of view of calculating what is its average displacement, it keeps moving. At the very beginning of the simulation, atoms move. And then they start feeling the potential from the other atom. And they sort of settle in an oscillatory movement.

And then you are starting to average over more and more oscillations. And you see sort of, in the bulk, you fairly quickly converge. These thin lines are for the bulk. You very quickly converge to your sort of final exact results.

But somewhere else for the surface atoms, it takes more time. It takes a few seconds before you finally converge to your sort of final result. But again, thermalization sort of can be checked by looking at the time evolution of the quantity that you want to average.

Say, you could look at the temperature of your system, and you want to see the quantity that you are sort of monitoring converge to a final result that doesn't change as you increase the length of your simulation. And there are quantities, again, that are sort of easy to converge and quantities that are very difficult to converge. Suppose for a moment that you want to study actually, again, the structure of liquid water.

Well, you could start from a sample that is frozen-- that is ice-- and all of a sudden you increase the temperature with your thermostat. You bring it, say, to 300 kelvin. And then you let it evolve.

And you could sort of look in time as the sort of relative position or the diffusion coefficient of the water molecule evolve in time. And you would discover that, depending on the temperature, maybe you need 10, 20, 30 picoseconds to reach a situation after which your expectation value do not change if you increase your simulation longer and longer. This is sort of the concept of equilibration.

If you were to do the opposite, if you wanted to find out what is the structure of ice by taking a sample of liquid water and cooling it down, your equilibration time would be almost infinite. Because if you take liquid water and you cool it down, it will actually not freeze in the perfect crystalline structure of ice. But it will much more likely sort of slow down in a sort of glassy transition in which it freezes.

It stop diffusing, but it's only sort of locally coordinated as ice should be. But it's not in a crystalline state. What does that mean?

It means that what we actually ended up with is a situation that is not at thermally representative situation. We are not in an equilibrium state for a system, but we have sort of and ended up in a corner of phase space out of which we are never going to get out unless we wait forever.

That is, even after their water is sort of frozen in this glassy state, if you keep simulating it longer and longer and longer, every now and then you'll have a fluctuation that changes the position of two water molecules and two more water molecules and two water molecules. And after a very long time, you might actually find that your system has become crystalline. But it's never going to happen in the timescale of your simulation.

So one needs to be careful when one does one of these simulation to make sure that, actually, you are sort of finding yourself in a representative state. And that's why some of these issues, like simulating a solid to liquid or a liquid to solid transition, require particular care and somehow require actually knowing something of the system that you want to discover.

And again, there is an enormous debate on what are some of the phases of supercooled water. That is, if you take water, and it's very pure, and you bring it down below the regular freezing temperature, well, you can keep water liquid up to, I think, 40 degrees below the freezing point. And you can do that experimentally. And you can do it with simulation.

And then, all of a sudden, your system will ultimately freeze. And it will end up in phases that neither the people doing simulation, nor the people doing experiment can still agree with. So there is sort of dark areas in the phase diagram of water at those low temperatures.

And people this discuss about high density liquid and low density liquid. And a lot has been achieved in the last 10 years. But it's still not formally known. And this is why one really needs to pay particular attention to a number of these problems.

OK. So suppose that everything is going well. You have initialized properly your system. You have sort of weighted enough, so that you have actually lost memory of your initial condition.

And if you want, you should always throw away the part of your molecular dynamic trajectory that corresponds to this initial thermalization time because it doesn't really represent high quality data. Only after you have waited long enough, say, in this previous simulation, only after you have waited, say, 4 picosecond, you can start taking averages of what is, say, the mean square displacement for a surface atom.

So you throw away from your simulation your thermalization time, and you start accumulating averages. And again, the averages that could be relevant, if you are seeing a microcanonical simulation, you could look at the average kinetic energy. Because that average kinetic energy gives you the average temperature of your system.

Or you could look at the average pressure in your system. Or I show you an example in a moment. You could look at the potential energy or what is sometimes called the caloric curve.

So if you are actually simulating a system at different temperatures-- say, again, suppose that you have water, and you are still in the frozen phase. Well, the potential energy will oscillate around a certain value. But remember, in order to sort of melt water-- that is in order to go at equilibrium from being an infinitesimal amount of temperature below the melting temperature and an infinitesimal amount above the melting temperature-- you need really to give to your system what is called the latent heat.

In this case, the latent heat of fusion would be the heat released when it cools or the heat that you give in order to melt it. And so your average potential energy will have a discontinuity as a function of temperature. And we'll see it in an example. And that's really a measure from your molecular dynamic simulation of this fundamental thermodynamical quantity.

Or you could look, say, in liquid water at the diffusion coefficient, so the means square displacements. And we'll see this. Or you could look at the structure, the geometric structure, of your liquid. And again, sort of molecular dynamics techniques were actually initially developed for system like liquids in which it was much more difficult to have analytical expression for something like the structure.

I mean, if you are studying a crystal, the structure is somehow trivial. And you can obtain it from X-ray experiments, but you can also obtain it from total energy calculation, minimizing the energy with respect to your volume. If you are studying a liquid, again, you can have information on your sort of local structure. That is how far the molecules keep on average from each other in the simulation from diffraction experiments.

But you don't have an immediate computational way to calculate data like you have done it for a solid in which you were, again, minimizing energy with respect to volume. What we really need is to do a molecular dynamic simulation. And this is sort of the first of the examples I want to show you today of application of molecular dynamics.

That is one in which we calculate ensemble average. And in particular, we, say, calculator, say, the structure of liquid water. I'll also give you an example, a couple of examples, of a chemical reaction that is using molecular dynamics to look at the evolution in real time of a system.

And I'll sort of hint at this more general idea that we can use these thermodynamical analogies and sort of molecular dynamic analogies to actually solve complex optimization problems. OK. And so this is an example from a molecular dynamic simulation, actually an ab initio molecular dynamic simulation, and one in which one tries to figure out what is the local structure of the liquid. And so how do you actually define structure in a liquid?

Well, sort of one of the fundamental quantities is what is called the pair correlation function. Suppose that we are looking, again, at water. We have the water molecules. So say we have an oxygen atom on each molecule.

And so what we can sort of ask ourselves is, what is the probability of finding another oxygen atom at a certain distance from one oxygen? So you could sort of imagine yourself sitting on a water molecule sitting on that oxygen and sort of monitoring on average what is the distribution of all the other oxygen atoms. And this would be one of these pair correlation function. This is specifically for the case of water.

I'll tell you in a moment what the difference between these curves is, but what we are plotting is the probability of finding an oxygen atom at a certain distance if I'm sort of sitting on one of these water molecules and running around. And as you can see, the probability below roughly 2.3, 2.4 angstrom is 0. So you never find an oxygen atom closer than 2.3, 2.4 angstrom.

And this is because, obviously, you know, molecules by Pauli repulsion and electrostatic repulsion, sort of repel each other very strongly at close distance. This, in the Lennard-Jones potential, was the $1/r^{12}$ term. When you get very close, the repulsion, it's so strong that you never get very, very, very close.

Actually, the only time you sort overcome this short range repulsion is in something like nuclear fusion in which you can manage to get, say, hydrogen atom to get so close together that actually their nuclei melt into a helium. But you see, at this standard condition of pressure and temperature, really the probability of finding some oxygen atom up to 2.3 angstrom is 0. And then all of a sudden, the probability of finding an oxygen atom skyrockets.

What does this mean? Well, it means that, actually, to have, if you add an oxygen atom in a water molecule, it's very likely to have other water molecules surrounding you at a distance that is a tad shy of 3 angstrom. OK. This is what the chemists would call the first solvation shell.

OK. The water molecule sort of runs around in the liquid. And there's a certain number of water molecules that are at a certain average distance that is around 2.6, 2.8 angstrom. And so this is a solvation shell, a shell of water molecules that, on average, surround this molecule.

And because there is a sort of average shell of water molecules around, once you move outwards from this oxygen, then there is going to be a depletion of water molecules. Because there is one shell at a certain distance. You move further, it becomes much more difficult to have water molecules at an intermediate distance.

And this is really sort of a general description for all kinds of liquids. And then you sort of move farther away. And then there is a second solvation shell.

OK, there is structure in the liquid. You have a water molecule. On average, it's surrounded by a first shell. And then, on average, that's surrounded by a second shell.

And as you move farther and farther, the boundaries become farther and farther. So you can sort of spot a third solvation shell. But once you move farther and farther away, if you look at what is the probability distribution 10 angstrom, 100 angstrom, 1 meter away, it becomes flat.

That is, if you look what is the probability of finding a water molecule 1 meter away or 1 meter plus 1/10 of an angstrom, there is really no difference. Because you have lost long range order. You have lost structure.

This would be very different, say, in a perfect crystal. So if you study something like silicon at 0 temperature-- and for a moment, forget about zero-point motion and quantum effects. Think just at a silicon crystal as being made of classical particles at 0 temperature. They are perfectly periodic.

And so what you would find, you could still define pair distribution function. And what you would find is Dirac's delta. That is the probability of finding a silicon atom at a certain distance from new silicon atoms, say, sitting on the origin is going to be 0 until you hit the distance of your first shell of neighbors.

And then it's sort of infinite or properly normalized at Dirac's delta. And then it's 0 again, and then another Dirac's delta, and 0 again, and then another Dirac's delta. And as you move farther and farther away, the density of your Dirac's delta becomes higher and higher.

Because, again, if you are at 1 meter of distance, there are an enormous amount of possibilities of finding sort of a combination of lattice vectors that give you a silicon atom in that position. And so if we were to smooth out those Dirac's delta, that is if we were to have some temperature, we would start seeing something a little bit like this. The first Dirac's delta would become slightly larger. The second Dirac's delta, the third Dirac's delta, we move farther and farther away.

There are so many close to each other that they all start overlapping. And also, for a crystal at final temperature, you would start to have this pair correlation function that goes towards 1. In a liquid, it's a much more dramatic.

And this quantity tells you also something fundamental about the appropriate size of your simulation cell. Remember that you are studying something in periodic boundary condition. So the question is, if we want to study liquid water, how large should our unit cell be?

I mean, does it need to contain 10^{23} , the Avogadro's number, of water molecules in a macroscopic sample? No. And this is a fundamental concept. It needs to be large enough so that the correlation between your molecule in the origin and that molecule periodically repeated is 0.

That is, by the time this pair correlation function reaches 1, it means that you have lost any structure. So somehow you don't have any way of recognizing that there is a molecule in the origin or that molecule is 0.1 angstrom displaced away. So a molecule that is on the origin and a molecule-- suppose that you continue this curve and maybe becomes 15 angstroms away. Well, those two molecules are not aware of each other's existence.

And so if you are studying your system, such data, the distance between a molecule and your periodic image is 15 angstrom. Actually, from all point of view, you are studying a system that behave as it was an infinite system. Your molecule sitting on the origin moves around.

And it sees independently all the molecules with which on average it interacts. And the molecule that is 15 angstroms away is so far away that there is no correlation. And so for all practical purposes, this molecule on the origin sees an infinity liquid system. It doesn't recognize that there is a boundary and things are periodically repeated.

And this is absolutely very important. And that's why we can actually simulate a realistic liquid with actually very few molecules. It turns out that, for something like water, you need only 32 water molecules at sort of regular conditions in the liquid, 300 kelvin, 350 kelvin, to actually simulate an infinite system.

You don't need to put more water molecules. It won't change anything. Of course, as you sort of get closer and closer to a phase transition, remember that, at the exact temperature of the phase transition, the solid and the liquid are in equilibrium.

So the free energy of the solid and the free energy of the liquid are identical. So it doesn't cost anything in terms of free energy to transform the solid in the liquid and the liquid in the solid. And that means that you can have fluctuation, bubbles of solid and bubble of liquid, that have arbitrary sizes.

So at the exact temperature of the phase transition, you have a divergence in this characteristic length. That is you can have fluctuations of all sizes. So if you were to study the system, exactly the phase transition, you need an infinite sized system.

But as you move away from the phase transition, very quickly the correlation function dropped down. And it's something like 25 kelvin above the melting point. It's so short that you need only 32 molecules.

And water is actually an excellent example to sort of show you also what goes wrong in the simulation. OK. Because what I've plotted here are actually a result of accurate or very accurate ab initio molecular dynamics simulation. And we know that sort of density functional theory actually describes very well the hydrogen bonds, so describes the interaction between water molecules very well.

But you can see it here. You see, these are the pair correlation function. You see this sort of set of curves here, the more structural set of curves, taken at 325, 350, 375 kelvin.

There is actually very, very little difference between them. And that's a bad sign. So what we are seeing is that this sort of structure is not changing as we increase the temperature.

And then we increase the temperature, say, from 375 to 400 kelvin. And you see this different set of curves. You see the red and the green.

So there is really a discontinuous change in going from 375 to 400 kelvin. What has happened in our simulation? Well, it turns out that, at 400 kelvin, our water is liquid.

At 375, it's a glassy system. So it's an amorphous glass. And you see it doesn't change with temperature. And you need to get to 400 kelvin for it to become liquid.

What is the bad thing here? It's not that the water is in an amorphous glass state. It's reasonable, as I said, even if it's not the result that we want, to find an amorphous states for a lot of liquids that have been cooled down below the melting temperature especially sort of liquids that tend to have sort of pseudo-covalent bonding.

You study carbon below its melting pot. You study silicon. You study something like water. They tend to want to be sort of tetrahedrally coordinated. And they tend to sort of form a lot of glassy structure.

What is really bad is that the melting of water doesn't seem to be taking place at 273 kelvin, but it takes place at 400 kelvin. OK. So if you want, if you were to use density functional theory and sort of molecular dynamics to predict what happens, we would all be frozen to death. You need to be at 400 kelvin to see something liquid. And you see, then you can also compare.

This is very good. There are a lot of experiments that look at what is the structure of water. You can do this with neutron scattering. You can do this with X-ray scattering. So you can probe the local structure. And you see this is what the experiment is.

So it's sort of good, but it's really not perfect. And so what's going on here? What's wrong in this?

Well, in many ways, we actually think that density functional theory is good enough to simulate water. OK. So the energy for any given configuration is given accurately by quantum mechanics. But there is something very important that I sort of hinted in some of the discussion.

300 kelvin, 400 kelvin is a temperature at which molecules can still display quantum effects in their vibrational degrees of freedom. That is when we discussed at certain point the fact that a simple molecule like hydrogen is actually only discrete vibrational energy and discrete vibrational states. And at 0 temperature, it's in the state called zero-point motion.

That is a purely quantum mechanical effect and doesn't really come from describing correctly the energy of the electrons in the potential of the ions, but comes from the fact that the ions themselves are quantum particles. And the excitation of a quantum particle from one state to another vibrational state is in itself quantized. And so what turns out to happen in water is that the hydrogen oxygen stretching frequency.

OK. So the frequencies by which one hydrogen beats against another oxygen are actually very high, are of the order of 3,000 centimeters to the minus 1. So that's a typical sort of measure of vibrational frequency. They are, in particular, sort of 10 times larger than the average kinetic temperature at room temperature.

So that means that, in reality, those vibrational modes cannot be excited. They are frozen in their ground state. So in all this liquid, you have a variety of vibrational modes, some that are very soft and that involve sort of how one water molecule correlates with other water molecules.

But you have some very stiff intermolecular vibrational degrees of freedom that you can't excite. And so the classical description of the ions is good, but it's not perfect. And that's what brings us this error here.

If we were to study-- sadly, there are no experiments, say, at 3,000 kelvin. But if you were to do an experiment of, obviously, water at a sort of fixed volume, so at very high pressure at 3,000 kelvin, then the ions would all sort of behave as classical particles and probably our agreements would be much better. But at this low temperature, some of these vibrational degrees of freedom are actually frozen.

And they can't exchange temperature. They can't exchange vibrational energy with the rest of the degrees of freedom. It's the same concept by which the specific heat of solids follows the Dulong and Petit law. And it's much smaller than its ideal sort of value at low temperatures.

Because even in a solid, where vibrational frequencies tend to be much lower than 3,000 centimeters to the minus 1, it tends to be of the order 300, 400 centimeters to the minus 1. So they tend to be all populated at room temperature. But in the regime of low temperatures, much lower 300 kelvin, some of those degrees of freedom are also frozen out.

So they can't exchange temperature when you increase the heat. And so the specific heat of the system is much smaller because a lot of those degrees of freedom can't actually talk with your heat bath. OK. So this was an example.

Let's see another example in which we look at ensemble averages. And this is the example of the caloric curve. So that is plotting the potential energy as a function of time.

This is done for a sort of classical molecular dynamic simulation. So you see it's using 6,400 particles, so a very large amount. This is actually a slab of lead in the 1, 0, 0 direction.

And you see, if we simulate this at 600 kelvin, well, this is as a function of time. This is a very long simulation, you see. We are talking about the nanoseconds.

OK. So it's of the order of 1 million time steps if we were using times step for a femtosecond. And you see, this is how the sort of average potential energy fluctuates at 607 kelvin. And then we increase the temperature, and you see it fluctuates around another value. Increase the temperature, it fluctuates another value.

We increase the temperature a little bit more from 625 kelvin to 630 kelvin, and what happens in this solid? You see, at the beginning, during, if you want, your thermalization time, you oscillate around a certain value. But then something happens, and your system goes into a new state in which the average potential energy oscillates around a new value.

What your system has done, it's gone through a melting transition. And so the difference between, if you want, the potential energy just below the melting temperature and the potential energy just above the melting temperature is this latent heat of fusion, the energy that is needed to melt the crystal or the energy that is released when the crystal solidifies.

And this is actually a sort of very specific system in which there is a sort of more exotic phase transition. And that's why molecular dynamics simulations were needed. It was a system in which there was what is called incomplete surface melting.

That is, if you look at the surface of lead 1, 0, 0, it starts melting at a temperature that is lower than the temperature at which the bulk melts. And this is because surface atoms are really different from bulk atoms. So they have their own phase transition. And it's temperature for that phase transition that is different from the bulk melting temperature.

So you heat up lead. All of a sudden, you see that the surface becomes liquid. And then you need to reach another critical temperature before also the bulk melts.

OK. So a couple of examples we have seen from ensemble averages. As I said, we can also use molecular dynamics to actually follow a chemical reaction. And so this is another example from ab initio molecular dynamic simulation.

Because what we really want to investigate is how we can take a molecule of methane-- you see C and H4. There are four hydrogen atom. One is sort of the pink and below, beneath the yellow. We want to figure out how to break this apart.

This is a fundamental technological problem because we have enormous reserves of methane is what people call natural gas. But it's extremely expensive to transport because you really need to liquefy it and carry it in sort of cryogenically frozen containers. So it's much more expensive to transport methane than to transport oil. And that's why we are actually not using sort of all the methane natural gas reserves that exist.

So one of the important things that we could do is actually transform methane in some higher and heavier hydrocarbons, like methanol, that we can transport much more easily. And that, in principle, is very doable. Methane has an enormous amount of energies ready to be released.

But sadly, in order to release the energy, in order if you want to burn methane or in order to break it apart, you need first to break apart the first carbon-hydrogen bond. And that's tough. Once you break that hydrogen bond, all the other hydrogen go and release a lot of energy.

But breaking this is very difficult. And we want to do this in a very controlled manner in order not to burn methane, but construct some higher hydrocarbon, like methanol. And so one sort of very suggestive hypotheses is actually using nanostructured catalyst, that is using well-known catalyst transition metals, like platinum in this case, but catalysts that have a very exotic surface, so that they become more reactive. They can help break part of this carbon-hydrogen bond.

And so this is where ab initio molecular dynamics come into play because we can study the system and try to figure out what happens when we take a methane molecule, and we throw it to a platinum surface. And as you can see, nothing happens. And we are very, very happy, not so much because we have sort of simulated an interesting system, but because actually we can measure the kinetic energy of the molecules coming down. And we can measure when it comes out.

And we find out that it has the same velocity, but opposite direction. So our system is conserving energy very accurately. But you know, what you can do now with your molecular dynamic simulation is try to figure out if this is the whole story.

That is, well, as you see, we have sort of thrown this molecule into the trough here. Because we thought that maybe this area is the most catalytically active. It's where there are more platinum atoms able to sort of rip this thing apart.

Well, it turns out that, actually, the most catalytically active part of the surface is the under-coordinated platinum atoms sitting here on the ridge. And you see, when the molecule falls down on this, it very easily breaks apart this. And so doing our ab initio MD, we discover that these under-coordinated atoms are actually much more active than the sort of surface-like atoms down here.

So under-coordinating a transition metal is actually very helpful. And this is actually what bacteria that eat methane for a living do. So what those bacterias do, they have an enzyme that is called a methane monooxygenase that actually is able to break apart this carbon-hydrogen bond. And what it has, it has sort of two iron atoms in the middle of a sort of porphyrin like structure. And those under-coordinated transition metals do all the trick.

Other examples in which a molecular dynamic simulation helps us understanding the structure of microscopic system-- as shown here. In this case, what we are doing is taking a liquid droplet of aluminum. So this is aluminum above its melting temperature. That is sort of 933 kelvin. It's been thermalized, and so it's a liquid.

And then we have a surface that could be, say, two different crystallographic orientation, the 1, 1, 0 surface and the 1, 1, 1 surface. And what we do is we drop this liquid droplet on the surface. And we sort of see what is the evolution of the system.

And on the 1, 1, 0 surface, that is one of these metal surfaces that undergo a pre-melting transition where the surface itself-- you can see it here-- becomes liquid-like. What happens is that you have a liquid droplet falling on a sort of liquid surface. And so it wets completely the surface, and this moves around.

You throw this same liquid droplet on aluminum 1, 1, 1, that is a surface that doesn't matter and that actually overheat. You can sort of bring a crystalite of aluminum, and you can look at the 1, 1, 1 surface. And you can bring it above the melting temperature, and it will not melt until you reach sort of a critical temperature of 1,000 kelvin.

Well, this liquid droplet actually doesn't wet away, but it forms, you see, a little disk. And you can actually sort of measure what is this contact angle. And so you can figure out what is the surface tension. And you can understand a lot basically about the free energy of the liquid and the solid just by those contact angles.

Another example for which I have another movie is sort of related to what you have seen in the sort of previous slide. And it's sort of looking at actually what do adatom on aluminum 1, 0, 0 surface do. And so you see what we do is, in ab initio molecular dynamics, we put an atom here. And we let the system evolve at a certain temperature.

In this case, it's evolving at 500 kelvin. That is roughly 1/2 of the melting temperature. And again, you can see some features of an ab initio molecular dynamic calculation in which we can simulate very few atoms.

And so actually, in this simulation, you see our surface area is given just by 16 atoms, this 4 by 4 here. And then we have periodic boundary condition. So this atom and this atom are doing exactly the same thing, but we are actually ourselves in the condition in which there is no correlation between what this atom does and what this atom does.

So these atoms sort of behave independently from what this atom does. Just because, at this temperature, the correlation length between these two is very, very short. And what I'm not showing here in this simulation is actually the lower layers of our aluminum slab just because they are not doing anything interesting.

And so we want to understand how the atoms move around. And everyone would have sort of guessed that adatoms on the surface hop around, sort of the classical model of diffusion, from one side here in the middle to another side here in the middle. But you see, you do your ab initio molecular dynamic simulation, and you discover that something very intriguing happens.

Actually, adatoms do not diffuse by hopping, but they diffuse by sort of diving into the surface layer and sort of pushing another atom out. And you'll see it again in a cleaner fashion. You see the diffusion takes place by an atom diving in.

Look at this. This sort of, in a moment, will dive in and pushes this sort of second atom out. And by the time it reaches here, now you see it appearing in periodic boundary condition.

And again, the molecular dynamic simulation has given you an insight on the microscopic process that provided density functional theory is accurate. And it, for this particular system, is actually very different from any kind of physical intuition would have told you. And it's 1 more minute.

Let me finish this movie because something more interesting takes place once you increase the temperature to 700 kelvin. This is what we have done in the simulation here. And what you actually see is that the first surface layer starts really behaving as a liquid. OK.

So we actually have a phase transition that is induced by the presence of that atom in which this first layer, if you monitor long enough what the mean square displacements are, how the atoms are, well, the atoms in this layer move around like a liquid. But they don't move around below the surface. And the surface remain crystalline.

And I'll finish the movie. We actually then, in our computational experiment, all of a sudden we can remove atoms from our liquid sample. That is we can sort of take them away. And we actually see very quickly that the system crystallizes again.

And you'll see it in a moment. See, we have removed atoms. We have actually removed two atoms. So we have created a vacancy. And the system has become crystalline.

Now, the atoms vibrate again around the crystalline surface. And there is a vacancy here. And from this, we also discover how the vacancy moves.

And you'll see that there is an atom from below that hops out. You see this hops out, leaving a hole here. And an atom from the first layer pops down.

So the vacancy that was here has moved here. And again, the way it has moved is not sort of the usual way in which one could imagine a vacancy moving by a surface atom moving around, but it actually involves a second layer atom jumping out.

OK. With this movie, I'll finish today's molecular dynamic simulation. We'll sort conclude the MD part on Thursday. And as a reminder, today was also the deadline for your second assignment, lab two. And you can give it to me or to one of the TAs. And otherwise, we'll meet on Thursday.