The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high quality educational resources for free. To make a donation or view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

PROFESSOR: And to this purpose, we need to calculate some thermodynamics. And we usually do that in statistical mechanics by calculating for some kind of a partition function. And we saw last time that it would be useful to calculate this grand partition function, which is the ensemble where you specify the temperature, the chemical potential, and the volume of the gas.

And in this ensemble, your task is to look at within this box of volume V the possibility of there being any number of particles. So you have to sum over all possible N particle states. The contribution of each N particle states is exponentially related to the number of particles through the chemical potential. And then given that you are in a segment that has N particles, you have to look at all possible configurations of those particles and integrate over all of those possibilities. And that amounts to calculating the partition function.

So for the partition function of an N particle system, you have to integrate over all of the momenta. Integration over each component of the momentum gives you a factor of 1 over lambda, where lambda again was related to the mass of the particle and temperature through this formula with h what we use to make these integrations over pq combinations dimensionless.

There are 3N such integrations. So that's a contribution from the momenta. And then we have to do the integration over all of the coordinates. And as long as these particles are identical, we decided to divide by the number of permutations. Because we cannot tear them apart.

So having done that, I need to now integrate over all of the N particles spanning a box of size V so the integration is within the box. And when I have interactions, then

I have to worry about the Boltzmann weight that comes from the interaction.

So here I should really put some kind of e to the minus beta times the interaction U. And what we did was we said that let's assume that this interaction comes from pairs of particles. And so this U is the sum over all possible pairs of particles, which when exponentiated will then give you a product over all possible pairs j, k and a factor that is related to the potential interaction between these.

And we found it useful to write that factor as 1 plus fjk, where this fjk stood for e to the minus beta of V qj minus qk minus 1. So basically if I add the 1 to that, I just get the exponentiated interaction potential. And then I have a forum such as this.

So basically, this is the quantity that we wanted to compute. And so then we said, well, let's imagine expanding these factors of 1 plus f1, 2, 1 plus f1, 3, 1 plus f2, 3, all of these factors, and organize them according to the powers of f that they have.

So the leading term would be taking 1 from everybody. So that would give me V to the N, which would be the 0-th order partition function, the non-interacting system. And then I would start to get corrections where there's order of f integrated, order of f squared integrated, all kinds of things.

And that being a somewhat difficult object to look at, we said, let's imagine graphically what we would get. And the typical contribution that we would get to this would involve having to iterate over all of these N particles. So we have to somehow imagine that we have particles 1 to N.

And then for a particular term, we either pick 1's-- and there will be some points that are not connected to any f that is integrated. There will be a bunch of things that will be connected to things where there are pair of f's. There will be things later on maybe where there are triplets of f, and so forth and so on.

And then we said that when I do the integrations that correspond to this, what do I get? I will get the contribution that comes from one particle by itself integrated. Let's call that b1 to the power of n1. I will get the contribution from this pair integrated. I will get b2 to the power of n2. I will get the contributions from these entities.

And in general, I said, well, OK, somewhere in this, I will get b_l to the power of n_l. Now of course, I have a big constraint here that is that the sum over l of ln_l has to be the total number of points 1 to n. So however I partition this, I will have for each graph that particular constraint acting.

We said that clearly there's a lot of graphs and combinations that give you precisely this same factor. But all I had to do was to sort of rearrange the numbers and ordering, et cetera, and I would get all of this. So it would be very nice if I could figure out what the overall factor is out here.

So we said that the factor is something like N factorial. Because what I can do is I can permute all of these numbers, and I would get exactly the same thing. But then I have to make sure that I don't over count. And not over counting required me to divide by the number of permutations that I have within each subgroup. So I have b_l to the power of n. I have l factorial to the power of n_l.

And then I have the change. Let's say this is 1, 2, this is 3, 4. I could have called one of them 3, 4, the other one 1, 2. So basically I will have n_l factorial from the permutations within each.

But I would have gotten exactly the same numerical factor out front if I had the same configuration but I had this diagram. I would have gotten exactly the same factor. I would call the contribution to this b3. I didn't say currently what b3 is. I would have gotten exactly the same factor.

So maybe then what I did was to sort of group all of those things that would come with the same numerical factor into this [INAUDIBLE] sum that I call b_l. So I call b_l to be the sum over all l-particle clusters. And then, of course here, I have to sum over all configurations of n_l that are consistent with this constraint that I have to put up there.

So the rest of it then was algebra. We said that if I constrain the total number, it's difficult for me to do. That's why I didn't go and calculate the partition function and switch to the grand partition function. Because in the grand partition, I can

essentially make this N that constrains the values of these nl's to be all over the place. And therefore summing over things with nl unconstrained is equivalent to summing over terms with nl constrained, and then summing over whatever the final constraint is.

So once I did that, I was liberated from this constraint. I could do the sum for each value of nl separately. And the thing nicely broke up into pieces. And so then what I could do is I could show that each term in the sum, there would be a product over different contributions l. For each l, I could sum over nl running from 0 to infinity.

I had the 1 over nl factorial from out here. I had an e to the beta mu divided by lambda cubed from the combination of these things raised to the power of lnl, which is how big N would have been composed. So I would have here lnl. And I also had an l factorial raised to the power of nl. So I can write things in this fashion.

So this thing then became the same thing as an exponential of a sum over l running from 1 to infinity. 1 over l factorial, of if you like e to the beta mu over lambda cubed to the power of l 1 over l factorial. And then I had bl. So it was this very nice result somehow summing over all kinds of things, and then taking the logarithm. The logarithm really depends only on the contributions of single clusters.

And again, the reason it had to be that way is because the ultimate thing that I calculated in this ensemble is that the answer should be e to the beta V times the pressure. And so the expression that we have over here better have terms which are all proportional to volume. Sorry, they're all made extensive by proportionality to volume.

And indeed, when I do the integrations over a single any cluster, there is one degree of freedom, if you like, associated with the center of mass of the cluster that can go and explore the entire volume. And so all of these things are in fact in the large end limit proportional to V and something that I call bl bar.

So once I divide by this volume, the final outcome of my calculation is that I can calculate the pressure of an interacting gas by summing over a series whose terms

are this e to the beta mu over lambda cubed raised to the power of l bl bar divided by l factorial.

OK, this is correct, but not particularly illuminating. Because the thing that we said we have some intuition for is that the pressure of a gas-- let's say if I look at it in terms of beta P, actually P is the density times kt, which is 1 over beta, if you like, or nkT. It doesn't matter.

And then there will be corrections. There will be a term that is order of n squared. There will be a term that is order of n cubed. And there are these coefficients, which are functions of temperature, that are called the virial coefficient. And this is a virial expansion.

Essentially what it is is a fitting of the form of the pressure of the gas as a function of density. In the very low density limit you get ideal gas result. And presumably because of these interactions, you will get corrections that presumably also know about the potential that went into the construction of all of these b bars, et cetera.

So how do we relate these things that are, say, experimentally accessible to this expansion that we have over here? And the reason it is not obvious immediately is because this is an expansion in chemical potential, whereas over here, I have density. So what do I do?

Well, I realize that the density can be obtained as follows. In the grand canonical ensemble, the number of particles is in principal a random variable. But that random variable is governed by this e to the beta mu N that controls how many you have.

And if I take the log of Q with respect to beta mu, I will generate the average number, which in the thermodynamic limit we expect to be the same thing as what we thermodynamically would call the number. And this log Q we just established is beta PV. So what we have here is the derivative of beta VP with respect to beta mu.

And if I do this at constant temperature, the betas disappear. This is the same thing as V dP by d mu at constant temperature. But also I can take the derivative right here.

So what happens? I will find that the density which is N over V, is the derivative of this expression with respect to beta mu. I go and I find that there's a beta mu. If I take a derivative of this with respect to beta mu, really it's exponential of beta mu l. The derivative of it will give me l e to the beta mu l.

So what I get is a sum over l e to the beta mu over lambda cubed. l gets repeated. And then I will have l times this. So I will have vl bar divided by l minus 1 factorial. So the series for density is very much like the series for pressure, except that you replace 1 over l factorial with 1 over l minus 1 factorial.

So now my task is clear. What I should do is I should solve given a particular density for what mu is. Once I have mu as a function of density, I can substitute it back here. And I will have pressure as a function of density.

Of course it's clear that the right variable to look at is not mu, but x, which is e to the beta mu over lambda cubed. This actually has a name. It's sometimes called fugacity. So the second equation is telling me that the density I can write in terms of the fugacity as a sum over l x to the l bl bar divided by l minus 1 factorial, which if I were to sort of write in its full details, it starts with b1 bar x plus b2 bar x squared. Because there I have 0 factorial or 1 factorial respectively.

The next term will be b3 bar over 2 x cubed. And this will go on and on. Also, let me remind you what b1 bar is. b1 bar, for every one of these b1 bars, I have to divide by 1 over V because of this V. And then I have to do the integration that corresponds to the one cluster, which is essentially one cluster going over the entirety of space, which will give me V. So this is in fact 1.

So basically this coefficient here is 1. And then I will have the corrections. So you say, well, I have n as a function of x. I want x as a function of n. And I say, OK, that's not that difficult. I will write that x equals n minus b2 bar x squared minus b3 bar x cubed over 2, and so forth. And I have now x.

Maybe a few of you are skeptical. Some of you don't seem to be bothered. OK, so my claim is that this is indeed a systematic way of solving a series in which when the

density goes to 0, you expect x to also go to 0. And to lowest order, x will be of the order of density. And these will be higher order powers in density.

So to get a systematic series in density, all you need to do is to sort of work with a series such as this keeping in mind what order you have solved things to. So I claim that to lowest order, this really just says that x is n. And then there are corrections that are order of n squared.

And to get the next order term, all I need to do is to substitute the lower order in this equation. So basically if I substitute x equals n in this equation, I will get n minus b2 n square. If I substitute x over n here, it will be a higher order term. And I claim that this is the correct result to order of n cubed.

And then to get the result to order of n cubed, I substitute this back into the original equation. So I start with n minus b2 bar, the square of the previous solution at the right order. So squaring this, I will get n squared minus twice b2 bar n cubed. And the next term that would be the square of this, which is order of n to the fourth, I don't write down.

And then I have minus b3 bar over 2, the cube of the previous solution. And at the right order that I have, that's the same thing as b3 bar over 2 n cubed and order of n to the fourth. And all I really need to do at this order is to recognize that I have two terms that are order of n cubed. Putting them together, I have n minus b2 bar n squared. And then I have plus 2b2 bar squared minus b3 bar over 2 n cubed and order of n to the fourth. So now we erase this.

OK, so now I have solved for x in a power series in n. All I need to do is to substitute in the power series for p. So let's write that down. The power series for beta p is b1 bar times the density b2 bar squared over 2 density square, b3 bar over 3 factorial, which is 6 n cubed, and so forth. But I only calculated things to order of n cubed. Also, b1 bar is the same thing as 1. So basically that's what I should be working with.

And so all I need to do is to substitute-- oops, except that these are all x's. The

series here for beta p that I have is in powers of this quantity x. And it is x that I had calculated as a function of n.

To the order that I calculated, it is n minus b2n squared. And then I have 2b2 squared minus b3 over b3. Somehow this sounds incorrect to me. Nope, that's fine. Because here I have 2. So that's correct. b3 over 2 n cubed-- so essentially, I just substituted for x here.

The next order term is going to be b2 over 2 times the square of x. The square of x will give me a term that is n squared, a term that is from here 2b2n cubed. Order of n to the fourth I haven't calculated correctly.

Order of n cubed I have b3 over 6. And just x cubed is the same thing as n cubed to this order. And I have not calculated anything at the next order. So let's see what we have.

We have n. At the next order, there are two terms that are n squared. There is this term, and there is that term. Putting them together, I will get minus b2 over 2 n squared. At the next order, at the order of n cubed, I have a bunch of terms.

First of all, there is this 2b2 squared. But then multiplying this with this will subtract one b2. So I'm going to be left with b2 squared. And then I have minus 1/2 b3. So that's one term.

And then I have plus 1/6 b3. Minus 1/2 plus 1/6 is minus 1/3. So this is minus b3 over 3 n cubed. And I haven't calculated order of n to the fourth. And this is the formula for BP.

So lowest order, I have the ideal gas result. Actually, let me, for simplicity, define the virial coefficients in this fashion.

The second order, I will get a correction B2, which is minus 1/2 of b2 bar. So this is minus 1/2 of the diagram that corresponds to essentially one of these lines that I have up here. And so what is it? It is minus 1/2 the integral over the relative coordinate of e to the minus beta v as a function of the relative coordinate minus 1.

Now earlier, we had in fact calculated this result directly through the partition function. I did a calculation in which I calculated the first term in the other way of looking at things, in the cumulant expansion, as a function of expansions in the potential. We saw that there was a term that all was order of density squared.

We summed all of the terms to get this factor. And there was precisely a factor of minus 1/2 as a correction to pressure once we took the derivative of the partition function with respect to volume. OK, so the thing that is new is really the next order term, B3, which is b2 squared minus b3 bar over 3.

OK, so what is that? Diagrammatically, B2 was this factor that we calculated above and is the square of this pair. And then I have to subtract from that 1/3 of whatever goes into B3. Now remember, we said that B3, I have to pick three points and make sure I link them in all possible ways. So the diagrams that go into B3, one of them is this.

But then there are three other ways of making a linked object, which is these things. We also saw that if I were to calculate the contribution of any one of these objects, I can very easily choose to measure my coordinates with respect to, say, the point that is at their apex. And then I would have one variable which is this distance, one variable which is that distance. And independently, each one of them would give me the factor that I calculated before.

In essence, all of these one particle reducible graphs give a contribution, which is the product of these single line graphs. There are three of them. Minus 1/3 precisely cancels that. And so the calculation that goes into the third virial coefficient ultimately will only depend on this one graph.

And we had anticipated this before. When we were doing things previously using the expansion of the partition function, we saw that in this cumulant expansion, only the cumulants were appearing. And for calculation of the cumulants, I had to do lots of subtractions. And those subtractions were genetically removing these one particle reducible graphs.

And this continues to all orders. And the general formula is that the l-th contribution here will be a factor of l minus 1 over l factorial, which is not so difficult to get by following the procedures that I have described over here, and then the part of bl that is one particle irreducible.

So this is the eventual result for how you would be able, given some particular form of the two body interaction, to calculate an expansion for pressure in powers of density, how the coefficients of that expansion are related to properties of this potential through this diagrammatic expansion.

So this is formally correct. And then the next question is, is it practical? Is it useful? So we need to start computing things to see something about the usefulness of this series for some particular type of potential.

So we are going to look at the kind of potential that I already described for you. That is, if I look at two particles in a gas that are separated by an amount r-- let's imagine the potential only depends on the relative distance, not orientation or anything else. We said that basically, at large distances, the potential is attractive because of van der Waals. At short distances, the potential is repulsive. And so you have a general form such as this.

Now, if I want to do calculations, it would be useful for me to have something that I can do calculations exactly with and get an estimate. So what I'm going to do is to replace that potential, essentially, with a hard wall. So my approximation to the potential is that my v of r is infinite for distances, separations that are less than some r0.

So basically, I define some kind of r0, which is the-- if you think of them as billiard balls, it's related to the closest distance of approach. This potential at large distances has typical van der Waals form, by which I mean it falls off as a function of separation with the sixth power. It is attractive. So I put a minus sign here.

In order to make sure that I get eventually the dimensions right, the coefficient that goes here I write as r0 with a u0 up here such that if I assume that the potential is

precisely this for all r that are greater than r0, then I'm replacing the actual potential by something like this. And the minimum depth of potential will occur at this cusp over here at minus u0.

So why did I do that? Because now I can calculate with that what the second virial coefficient is. So what is B2? B2 is minus 1/2 the integral of relative potential. Now this integral will take two parts from these two different contributions. One part is when I'm in the regime where the particles are excluded, the potential is infinity. And there, f is simply minus 1.

So that minus 1 then gets integrated from 0 to r, giving me the volume that is excluded. So what I will get here is a minus 1 times the excluded volume. That's the first part, where the volume is, of course, 4 pi over 3 r0 cubed for the volume of the sphere.

And then I have to add the part where I go from f0 all the way to infinity. The potential is vertically symmetric. So d cubed r becomes 4 pi r squared dr. And I have to integrate e to the minus beta times the attractive part of the potential. So it's beta u0 r0 over r to the sixth minus 1.

So what I will do, I will additionally assume that I'm in the range of parameters where this beta u0 is much less than 1, that is, at temperatures that are higher compared to the depth of this potential converted to units of kT. And if that is the case, then I can expand this. And the expansion to the lowest order will simply give me beta u0 r0 over r to the sixth power. And I will ignore higher order terms. So this, if you like, is order of beta u0 squared.

Now, having done that, then the second integral becomes simple. Because I have to integrate r squared divided by r to the fourth, which is r to the minus 4. Integral of r to the minus 4 gives me r to the minus 3. And then there's a factor of minus 1/3 from the integration. It has to be evaluated between 0 and infinity.

And so the final answer, then, for B2 is minus 1/2. I have minus 4 pi over 3 r0 cubed from the excluded volume part. From here, I will get a 4 pi beta u0 r0 to the sixth.

11

And then I will get this factor of r the minus 3 over 3 evaluated between infinity, which gives me 0 and r0.

So you can see that with this potential, actually both terms are proportional to 4 pi r0 cubed over 3. And so I can write the answer as minus omega over 2 where I have defined omega to be this 4 pi r0 cubed over 3, which is the excluded volume. Basically, it's the volume that is excluded from exploration when you have two particles. Because their center of mass cannot come as close as r0.

So that 1/2 appears here as 1. Because I already took care of that. This rest of the coefficients are proportional to this except with a factor of minus beta u0. So I get minus beta u0. So this is your second virial coefficient for this.

AUDIENCE: It should be positive out front.

PROFESSOR: Pardon?

AUDIENCE: It should be positive out front, [INAUDIBLE]?

PROFESSOR: Yes, because there was a minus, and there's a minus. There's a plus. Let's keep that. So what we have so far, we've said that beta times the pressure starts with ideal gas behavior n. And then I have the next correction, which is this B2 multiplying n squared. So I will have n squared times this coefficient that I have calculated for this type of gas, which is omega over 2 1 minus beta u0. And presumably there are higher order terms in this.

So whenever you see something like this, then you have to start thinking about, is this a good expansion? So let's think about how this expansion could have become problematic. Actually, there is already one thing that I should have noted and I didn't, which is short versus long range potential.

Now, it was nice for this potential that I got an answer that was in the form that I could write down, factor out omega. You can say, well, would you have always been able to do something similar to that? Because you see, ultimately what this says is that in order for all of these terms in this expansion to be dimensionally correct, the

next term is a density, has dimensions of inverse volume, density squared compared to this. So it should be compensated by some factor of volume.

And we can see that this factor of volume came from something that was of the order of this size of the molecule, this r0. OK, so it's interesting. It says that it's really the short-range part of the potential that seems to be setting the correction.

Now, where did that come from? Well, there was one place that I had to do an integration over the potential. And I found that the integration was dominated by the lower range. That was over here. Where would this become difficult?

AUDIENCE: In a small box?

PROFESSOR: In a small box, but we are always taking the thermodynamic limit where V goes to infinity.

AUDIENCE: Or inverse square log?

PROFESSOR: Inverse square log, yes, that's right. So who says that this integral should converge? I'm doing an integral d cubed r, something like v of r. And convergence was the reason why this integral was dominated by short distance. If this potential goes like 1 over r cubed, then it will logarithmically depend on the size of the box. For Coulomb interaction, 1/r potential you can't even think about it. It's just too divergent.

So this expansion will fail for potentials that have tails that are decaying as 1 over r cubed or even slower. Fortunately, that's not the case for van der Waals' potential and typical potentials. But if you have a plasma, you have to worry about this. And that's why I was also saying last time that the typical expansions that you have to do for plasmas are different.

So given that we are dominated by the short range, what is the correction that I've obtained compared to the first term? So essentially, I have calculated a second order term that is of the order of this divided by the first order term and pressure that was density. And we find that this is of the order of n omega.

How many particles are within the range of interaction? And this would be kind of like the ratio of the density of liquid to the density of gas. Because the density of liquid would be related-- oh, it's the other way around. The density of liquid would be 1 over the volume that one particle occupies. So the density of liquid would go in the denominator. The density of gas is this n that I have over here.

And again, for the gas in this room, this ratio is of the order of 10 to the minus 3. And we are safe. But if I were to start compressing this so that I go to higher and higher densities, ultimately I say that this second order term becomes of the order of the first order term. And then perturbation doesn't make sense.

Naturally, the reason for that is I haven't calculated. But typically what you find is that if you go to higher and higher order terms in the series, in most cases, but clearly not in all cases, the ratio of successive terms is more or less set by the ratio of the first terms. It has to be dimensionally correct. And we've established that the typical dimension that is controlling everything is the volume of the particle.

So this expansion will fail for long range potentials, for going to liquid-like densities. But also I made another thing, which is that I assumed that that could expand this exponential. And then in principle, there are other difficulties that could arise if this condition that I wrote here is violated. If beta u0 is greater than 1, then this coefficient by itself becomes large exponentially.

And then you would expect that higher order terms will also get more factors of this exponential and potential. And things will blow up on you. So it will also have difficulties at low temperatures for attractive potentials.

Again, the reason for that is obvious. If you have an attractive potential, you go to low enough temperature, and the ground state is everybody sticking together-- looks nothing like a gas, looks like a solid or something. So these are the kind of limitations that one has in this series.

OK, let's be brave and do some more rearrangements of this equation. So I have that to this order, beta P is n plus n squared this excluded volume over 2 1 minus

beta u0. And then there's higher order terms of course. And then I notice that there are two terms on this equation that are proportional to beta. And I say, why not put both of them together?

So I will have beta. Bring that term to this side, and it becomes P plus n squared omega over 2 u0. And then what is left on the other side? Let's factor out the n. I expect this to be a series in higher and higher powers of n.

And the first correction to 1 comes from here, which is n omega over 2. And I expect there to be higher order terms. Now again, to order of n squared that I have calculated things correctly, this expression is no different from the following expression-- 1 minus n omega over 2. Again, there will be higher order terms in n. But the order of n squared, both of these equations, expressions, are equivalent.

So if I now ignore higher order terms, this whole thing is equivalent to P plus n squared u-- well, let's write this in the following fashion-- p plus u0 omega over 2 N over V, which is density squared. And actually this side, if I were to multiply by V, what do I get? This becomes numerator V, denominator V minus N omega over 2. Let's multiply by V minus N omega over 2. And the right-hand side will be--

AUDIENCE: Is that where it should be there?

PROFESSOR: What happened?

AUDIENCE: This one, [INAUDIBLE].

PROFESSOR: Yes, so if I multiply this, which is N over V, by V, I will get N, good, which removes the difficulty that I had. Because now I multiply by kT. And the left hand side disappears. And on the right hand side, I will get NkT.

So we'll spend some time on this equation that you will likely recognize as the van der Waals equation. I kind of justify it by rearranging this series. But van der Waals himself, of course, had a different way of justifying it, which is that basically if we think about the ideal gas, and you have particles that are moving within some volume V, if you have excluded volume interactions, then some of this volume is no

longer available. And so maybe what you should do is you should reduce the volume by an amount that is proportional to the number of particles.

This factor of 1/2 is actually very interesting. Because it is correct. And I see that a number of people in well known journals, et cetera, write that the excluded volume should be essentially n times the volume that is excluded around each particle.

It's not that. It is 1/2 of that. And I'll leave you to mull on that. Because we will justify it later on. But in the meantime, you can think about why the factor of 1/2 is there.

The other issue is that-- so there has to be a correction to the volume. And the kind of hand-waving statement that you make about the correction to pressure is that if you think about the particle that is in the middle, it is being attracted by everybody, whereas when it comes to the surface, it is really being attracted by things that are half of the space.

So there is an effective potential that the particles feel from the collective action of all the others, which is slightly less steep when you approach the boundaries. And therefore, you can either think that because of this, there's less density that you have at the boundary. Less density will give you less pressure.

Or if you have a particle that is kind of moving towards the wall, it is being pulled back so it doesn't hit the wall as strongly as you would expect, that would give the pressure of the ideal gas. So there is a pressure that has to be reduced related to the strength of the potential and something that has to do with all of the other particles. And there's density squared will appear there.

We will have a more full justification of this equation later on. But for the time being, let's sort of sit with this equation and think about its consequences for awhile. Because the thing that we would like to do is we have come from a perspective of looking at the ideal gas and how the pressure of the ideal gas starts to get corrected because of the interactions.

Of course, things become interesting when you go to the dense limit. And then the gas becomes something like a liquid. And you have transitions and things like that.

So really, it's the other limit, the dense, highly interacting limited that is interesting.

And to get that, we have few choices. Either I have to somehow sum many, many terms in the series, which will be very difficult, and we can't do that, or you can make some kind of approximation, rearrangement, and a guess. And this is what the van der Waals equation is based on.

I made the guess here by somewhat rearranging and re-summing the terms in this series. But I will give you shortly a different justification that is more transparent and tells you immediately what the limitations are. But basically that's why we are going to spend some time with this equation. Because ultimately, we are hoping to transition from the weakly interacting case to the strongly interacting case. Yes.

AUDIENCE: [INAUDIBLE]

PROFESSOR: No, no, no, no, omega was defined as the volume excluded around one particle. So if you're thinking about billiard balls, r0 is the diameter. It's not the radius. And 4 pi over 3 r0 cubed is 8 times the diameter of a single billiard ball. So the correction that we get, if you like, is 4 times the volume of a billiard ball multiplied by the number of billiard balls. Where was I?

OK, so what this equation gives you, the van der Waals, is an expression for how the pressure behaves as a function of volume. Actually, it would be nicer if we were to sort of replace this by volume per particle, which would be the inverse density. But you can use one or the other. It doesn't matter.

Now, what you find is that there is, first of all, a limitation to the volume. So basically, none of your cares are going to go to volumes that are lower than n omega over 2. So basically there's a barrier here that occurs over here.

But if you go to the other limit, where you go to large volumes, you can ignore terms like this. And then you get back the kind of ideal gas behavior. So basically, in one limit, where you are either at high temperatures-- and at high temperatures, essentially, the correction here will also be negligible. Or you are at high values of the volume. You get isotherms that are very much like the isotherms that you have

17

for the ideal gas, except that rather than asymptote to 0, PV going like NkT, you asymptote to this excluded volume.

So this is for high T, T large. T larger than what? Well essentially, what happens is that if I look at the pressure, I can write it in this fashion also. It is NkT divided by V minus N omega over 2. So that's the term that dominates at high temperature. It's proportional to kT.

But then there's the subtraction u0 omega over 2 N/V squared. And again, this term is not so important at large volume. Because at large volume, this 1/V is more dominant than 1 over V squared. But as you go to lower temperatures and intermediate volumes, then essentially you have potentially a correction that falls off as minus 1 over V squared.

And so this correction that falls off as 1 over V squared can potentially modify your curve, bring it down, and then give it a structure such as this. So this is T less than. And clearly, between these two types of behavior, where there is monotonic behavior or non-monotonic behavior, there has to be a limiting curve that, let's say, does something like this, comes tangentially to the horizontal axis, and then goes on like this.

So this would be for T equals Tc. This is T greater than Tc, T less than Tc. OK, so that's fine, except that now we have encountered the difficulty. Because one of the things that we had established for thermodynamic stability was that delta P delta V had to be negative. And the ideal gas curve and portions of this curve which have a negative slope are all consistent with this. But this portion over here where dP by dV is positive, it kind of violates the condition that the compressibility kappa T, which is minus 1/V dV by dP at constant temperature, better be positive.

And so clearly, the expression that one gets through this van der Waals equation has a limitation. And the most natural way about it is to say, well, the question that you wrote down is clearly incorrect. That's certainly true.

Also, this equation is incorrect, except that this kind of reminds us with what actually

happens if I look at the isotherms of a gas, such as gas in this room, or something that is more familiar, such as water. What you find is that at high temperatures, you indeed have curves that look like this.

But at low temperatures, you liquidify. And what you have is a zone of coexistence, let's say something like this, by which I mean that the isotherm that you draw has a portion that lives in the gas phase and is a slightly modified version of the ideal gas. But then it has a portion that corresponds to the more or less incompressible liquid. Although this has clearly still some finite compressibility.

And then in between, there is a region where if you have a box, part of your box would be liquid, and part of your box would be gas. And as you compress the box, the proportion of liquid and gas will change. And this happens for T less than some Tc. And once more, there is a curve that kind of looks like this at some intermediate temperature Tc.

So you look at the comparison, you say, well, as long as I stay above Tc-- of course I don't expect this equation to give the right numbers for what Tc or whatever is. But qualitatively, I get topologies and behaviors at that T greater than Tc all the way up to Tc are not very different. But at T less than Tc, they also signal that something bad is happening.

And somehow, the original description has to be modified. And so the thing that Maxwell and van der Waals and company did was to somehow convert this incorrect set of equations to something that resembles this.

So let's see how they managed to do this. Actually, this may not be a bad thing to keep in mind. Let's have a thermodynamic-- well, I kind of emphasize that there is thermodynamic reason for why this is not valid. Always there's a corresponding reason if you look at things from the perspective of statistical mechanics. So it may be useful to sort of think about what's happening from that perspective.

Let's imagine that we are in this grand canonical ensemble. In the grand canonical ensemble, the number of particles, as we discussed, is not fixed. But the mean

number of particles is given by the expressions that we saw over here, is related to the pressure through dP by d mu at constant temperature.

But there are fluctuations. So you would say that the fluctuations are related by taking a second derivative of this object. Because clearly, q, if I were to expand it in powers of beta mu, will generate various moments of n. Log of q will generate cumulants. So the variance would be one more derivative.

And so that will amount to taking a derivative of the first derivative, which was the number itself. So I will get this to be dN by d mu at constant T, except that there's an additional factor of kT. So maybe I will write this carefully enough.

I have d2 log Q with respect to beta mu squared will give me N squared. Now, the first derivative with respect to beta mu of log Q gave me the mean, which I'm actually thinking of as N itself. And then there's this factor of beta. So this is really kT dN by d mu. All of these are done at constant temperature.

So the statistical analog of stability really comes down to variances being positive. So my claim is that this variance being positive is related to this stability condition. Let me do that in the following fashion. I divide these two expressions. I will get the average of N squared divided by average of N, which is really N, so the variance over N. I have kT. And then I have the ratio of two of these derivatives.

And the reason I wanted to do that ratio was to get rid of the d mu. Because d mu is not so nice. If I take the ratio of those two derivatives, I will get dN by dP at constant T.

But what I really wanted was to say something about dV by dP at constant T rather than dN by dP. So what I'm going to do is to somehow convert this dN by dP into dN by dV, and then dV by dP. And actually to do that, I need to use the chain rule. So this will be T, P. This will be T, N. And the chain rule will give me an additional minus sign.

So this is a reminder of how your partial derivatives and the chain rule goes. Finally, dN by dV is none other than the density N/V. At constant temperature, N and V will

be proportional. And what I have here is 1/V dV by dP, which is the compressibility. So the whole thing here is NkT times kappa of T.

And so this being positive, this positivity of the variance, is-- the thermodynamic analog of it was from the stability, and something such as this. The statistical analog of it is that the probability distribution that I'm looking at in the grand canonical ensemble as a function of N has to be peaked around some region.

The variance around that peak better be positive so that I'm looking in the vicinity of a maximum. If it was negative, it means that I'm looking in the vicinity of a minimum. And I'm looking at the least likely configuration. So it cannot be allowed. So that's the statistical reason.

OK, now I want to use that in connection with what we have over here and see what I can learn about this. I said that I wrote two expressions and divided them, because I didn't want to deal with the chemical potential. But let me redefine that and say the following, that N is V dP by d mu at constant T. And maybe I will start to gain some idea about chemical potential if I rearrange this as d mu is V/N dP along this surface of constant P at constant T. And then I can calculate mu as a function of P at some particular temperature minus mu at some reference point by integrating from the reference point to the pressure of interest the quantity V/N dP prime where V of P prime I take from this curve.

So this curve gives me P as a function of V. But I can invert it and think of V as a function of P, put it here, and see what is happening with the chemical potential if I walk along one of these trajectories. And in particular, let's kind of draw one of these curves that I am unhappy with that have this kind of form in general and calculate what happens to the chemical potential if I, let's say, pick this reference point A, and the corresponding pressure PA, and go along this curve, which corresponds to some temperature T that is less than this instability temperature, and track the shape of the chemical potential.

So this formula says that the change in chemical potential we obtained by calculating V-- divide by N, but our N is fixed-- as a function of P, and integrating as

you go up in pressure. So essentially, I'm calculating the integral starting from point A as I go under this curve.

So let's plot as a result of that integration how that chemical potential at P minus PA is going to look like as I go in P beyond PA. So at PA, that's my reference. The chemical potential is some value. As I go up and up, because of this area of this curve, I keep adding to the chemical potential until I reach the maximum here at this point C. So there is some curve that goes from this A to some point C.

Now, the thing is that when I continue going along this curve down this potential all the way to the bottom of this, which I will call D, my DP's are negative. So I start subtracting from what I had before. And so then the curve starts to go down. So all the way to D, I'm proceeding in the opposite direction.

Once I hit D, I start going all the way to the end of the curve however far I want to go. And I'm adding some positive area. So the next part I also have a variation in chemical potential that goes like this.

So you ask, well, if I give you what the temperature is and what the pressure is-- and I have told you what the temperature and pressure are-- I should be able to calculate the chemical potential. It's an intensive function of the other two intensive functions. And this curve tells me that, except that there are regions where I don't know which one to pick. There are regions where it's obvious there's one value of the chemical potential. But there is in between, because of this instability, three possible values.

And typically, if you have systems-- think about different chemicals that can go between each other. And there are potentials for chemical transformations. You say, go and pick the lowest chemical potential. The system will evolve onto the condition that has the lowest chemical potential.

So that says that when you have a situation such as this, you really have to pick the lowest one. And so you have to bail out on your first curve at the point B, and on your second curve at a point E, where on this curve B and E are such that when you

integrate this BDP all the way from B to E, you will get 0, which means that you have to find points B and E such that the integral here is the same as the integral here. And this is the so-called Maxwell construction.

So what that means is that originally I was telling you over here that there is some portion of the curve that violates all kinds of thermodynamics. And you should not access it. And this other argument that we are pursuing here says that actually the regions that you cannot access are beyond that.

There's a portion that extends in this direction, and a portion that extends in that direction. And really there's a portion that you can access up to here, and a portion that you can access up to here. And maybe you should sort of then join them by a straight line and make an analog with what we have over there.

Now, in another thing, another way of looking at this is that the green portion is certainly unstable, violates all kinds of thermodynamic conditions. And people argue that these other portions that have the right slope are in some sense metastable. So this would be a stable equilibrium. This would be an unstable equilibrium. And a metastable equilibrium is that you are at a minimum, but there's a deeper minimum somewhere else.

And the picture is that if you manage to take your gas and put additional pressure on it, there is a region where there is a coexistence with the liquid which is better in terms of free energy, et cetera. But there could be some kind of a kinetic barrier, like nucleation or whatever, that prevents you to go there.

And this again is something that is experimentally observed. If you try to rapidly pressurize a gas, you could sometimes avoid and not make your transition to the liquid state at the right pressure. But if you were to do things sufficiently slowly, you would ultimately reach this point that corresponds to the true equilibrium pressure of making this transition.

Now, the thing is that somehow this whole story is not very satisfactory. I started with some rearrangement of some perturbative to arrive at this equation. This

equation has this unstable portion. And somehow I'm trying to relate something that makes absolutely no thermodynamic sense and try to gain from it some idea about what is happening in a real liquid gas system.

So it would be good if all of this could be more formally described within a framework where all of the approximations, et cetera, are clear, and we know what's going on. And so next time, we will do that. We will essentially developed a much more systematic formalism for calculating the properties of an interacting gas, and see that we get this, and we also understand why it fails.