

[SQUEAKING]

[RUSTLING]

[CLICKING]

**GIAN PAOLO
BERETTA:**

So good morning. We're still playing with mixtures, properties of multi-component systems. We've defined some ideal behavior, like the Gibbs-Dalton mixture that applies for ideal gases. And today, we'll extend that to the case of liquid mixtures.

Remember, this was last-- the last slide of the last lecture. And we were discussing this Lennard-Jones interactions between particles. And remember that in the gas phase, when we can apply the ideal gas model, the particles, on average, live far away from each other. They do experience the collision very often.

And so very often, they do experience this acceleration towards the other particle and then the repulsion that pushes it back and then finishes the collision. But the time that they spend in this area is much smaller than the time they spend far away from each other. And therefore, for the calculation of the average properties at stable equilibrium, each particle behaves as if it were by itself. And in fact, that's the behavior that we have observed also for mixtures.

And so also for mixtures, whether the particles are close to similar particles or to different particles, if we are in the Gibbs-Dalton mixture limit, the same holds also for different particles. But you can see that if there are mixtures in which, for example, the individual components, when pure, stay far from each other, but when mixed like to stay close, in that case, you would have a non-ideal behavior because if they spend more time in this close-by area, then this will influence. And you cannot neglect anymore this energy of interaction.

Now, similar to the Gibbs-Dalton mixture behavior, we can define also an ideal solution behavior for the liquid phase, like sketched here. So in the liquid, of course, the particles live very close to one another. And so they live very close to this dip in the Lennard-Jones potential.

But we can consider an ideal situation-- hypothetical, of course. But it's a good start for any kind of modeling, or it could be also taken-- it's actually taken, typically, as a reference. And then we will describe actual real solution behavior as departures with respect to the ideal solution behavior.

This ideal solution behavior is one in which the-- this effect of the energy of interaction is alike for the particles for the-- for-- say, for the A particles and for the B particles. And it's also alike for the A and B particles. So essentially, these particles are very similar to one another. And they don't really care if the neighbor is similar or not.

In that limit, therefore, there is no contribution to the enthalpy of mixing because you remember, the enthalpy of mixing has to do with the fact that, for example, here, if I take a gray particle and I put it in a container full of black particles, the enthalpy of mixing is the change in the energy if I keep at the same temperature and pressure-- due to mixing-- adding 1 mole of that gray substance into the black bulk. And so if the particle lives in this area, the enthalpy of mixing will be related to this dip in the energy.

Now, this other case-- for the liquid case, we are somewhere here. So the energy is somewhere here. But it's the same for all. And so even if you add one more gray particle, it doesn't really make a difference for the energy. It doesn't really make a difference for the volume, also, because being at the same-- leaving it the same distance, the volume also doesn't change for the mixture.

But it does make sense-- I mean, it should make a difference regarding the entropy because as you remember, when we did this entropy of mixing idea, the fact that you have added one more gray particle-- you have diluted the mixture. So you have changed, say, what we call the surprise when you see a gray molecule. Now you've changed the probability of seeing the gray molecule with respect to the black one. And that is accounted for by the entropy of mixing. That's why the entropy of mixing does-- is the only effect. So in this ideal solution behavior, we only have this entropy of mixing effect.

Now, if you remember, the chemical potential for any constituent was written this way as the chemical potential of the pure component at the same temperature and pressure of the mixture plus the enthalpy of mixing of that component minus the temperature times the entropy of mixing. And we've used this relation already a few times.

So if I substitute the three assumptions up here-- so let me do first the volume. For the volume, I have to remember that the volume of the mixture is given, as we have proven by-- in general, as a weighted sum of the partial volumes of the constituents in the mixture. And the partial volumes can be written as the specific volume of the pure component at the same temperature and pressure of the mixture plus the volume of mixing.

So in this case, the volume of mixing is assumed to be 0 for ideal solution behavior. And so all we are left is that the volume is given by the weighted sum of the volumes that the single pure components would occupy at the same temperature and pressure. And this is Amagat law of additive volumes, something that we proved is satisfied also by a Gibbs-Dalton mixture of ideal gases.

For the chemical potential-- so Δx_{mix} -- here, there should be an equal 0. I thought I corrected this. So equals 0-- and equals 0. So Δx -- Δh is equal to 0. And Δs is equal to minus $R \log$ of y , which is the surprise multiplied by the Boltzmann constant, or the gas constant. So the minus and the minus become plus. And so the chemical potential becomes plus $RT \log$ of y .

At this stage, ideal solution behavior is when this happens for all constituents. And this is a very rare practical situation. But again, it's a good starting point, a good reference situation on which you can make some calculations.

Now, this ideal mixture solution behavior can be used to derive an important aspect of equilibria across semipermeable membranes, which is the osmotic pressure concept. And osmotic pressure is nothing but the difference between the pressure of the mixture and the pressure that the pure component has when in mutual equilibrium on the other side of a membrane that is permeable only to that component-- so it's p minus p_{ii} .

Now, the expression that I just reminded us on the other slide just behind this is rewritten here, by remembering that, of course, across that membrane, there is equilibrium, mutual equilibrium. And the condition for mutual equilibrium is that the chemical potential of the constituent in the mixture is equal to the chemical potential of the pure constituent on the other side of the membrane-- so up here, where the pressure is, the partial pressure.

And also here, I see another error to be corrected on the slides. This should be double i here and double i here. So this is equal to-- this is the chemical potential-- this is equal to the chemical potential of the constituent in the mixture. But it is equal to the pure chemical potential at the partial pressure.

So as you remember, when we calculate this difference, since we know that the derivative of the chemical potential with respect to pressure is the specific volume, this integral can be easily done once we know the equation of state of the substance, of the pure substance.

So this is nothing new. But it's the starting point of this observation made by van 't Hoff-- namely, that when we have a solution and we have a dilute solution-- so a dilute solution means that there is a lot of a substance. Think of, say, a lot of gray particles and a few solutes.

One is called the solvent. The others are called the solutes. And the solvent is the predominant amount. So that means that the mole fraction of the solvent is very close to 1. And 1 minus the mole fraction of the solvent, which is equal to the sum of the mole fractions of the solutes, is much less than 1. So that is a dilute solution.

Now, in a dilute solution, if you put in a gray particle in the box in which there are already many gray particles, it's very likely that that particle will end up with neighbors that are gray. So in this situation, there is no enthalpy of mixing. There is no change in volume of mixing. So it's exactly the kind of behavior that describes the ideal solution behavior for the solvent only, though.

So what we are saying is that of-- van 't Hoff's observations is that in any dilute solution, even if we are not in that limiting situation in which the blue-- I mean, the dark and the gray particles are identical in terms of intermolecular interactions, even if they're very different-- but the fact that there are few black particles makes that the addition of a gray one doesn't change the energy and the volume of the mixture. So in other words, the solvent of a dilute solution behaves according to the ideal solution model.

And of course, that doesn't mean that there is no effect in adding a gray particle because, again, if we add a gray particle, it is going to dilute further the solution. And that will have an effect on the entropy of mixing.

For example, if you take the entropy of mixing here, which is that expression that we derived, and you write it for this solution here-- the reason I want to write it carefully is that I want to compute the entropy of mixing. And that is the derivative of S not with respect to the mole fractions, but with respect to the amount of constituent i , 1, which is the solvent.

And so I have to write explicitly how the mole fractions depend on the amount. And you see there is an n_1 in the numerator. But there is also an n_1 in the denominator. So when you take the derivative, you have to be careful. Well, if you do it, it turns out that the change in this entropy of mixing by-- due to the addition of a small amount of n_1 is, indeed, equal to minus $R \log$ of y . And that is the surprise.

So this backs up van 't Hoff's observation. So at least for the solvent, we can use the ideal mixture solution. So we can set Δh equal to 0 here. And Δs_{mix} equals to minus $R \log$ of y . So this becomes a plus $RT \log$ of y . And here it is.

And for the other side, you see this integral. Now, if we apply it to a liquid-- so if we are talking about a dilute liquid solution, the specific volume of a liquid under the incompressible ideal fluid is constant. It doesn't change with temperature and pressure. So this v goes out of the integral. And it becomes the specific volume of the solvent times the difference in the integration limits. And that's the osmotic pressure.

So we got the expression here for the osmotic pressure. Actually, van 't Hoff makes an additional approximation, noticing that since the assumption is that the solution is dilute, it means that the mole fraction of solute is very close to 1.

So we can write it as 1 minus the mole fractions of the solutes. And then we can take the first term in the Taylor series expansion of \log of 1 minus x , which is minus x . So this becomes simply the summation of the solutes.

And finally, we have this expression, which is the van 't Hoff relation-- says the osmotic pressure of the solvent is equal to RT divided by the specific volume. And again, here, this should be a v_{11} , another error, times the sum of the mole fractions of the solutes.

Notice that it doesn't really matter what the solutes are. It can be any molecule. It could be sugar, could be chlorine, and so on-- could be salt, except that, as we will see in a moment because that is the important application that we want to discuss, as you know, some molecules-- for example, salt-- when you put it in water, they separate. They dissociate. And then they get solvated. They hydrolyze.

And so you have to be careful when counting the molecules because when you have, for example, the mole-- you put in a certain fraction of salt and then you-- it gets dissolved, for each mole that you put in, you get 2 moles of solutes. So you have to double this.

Now, clearly, this issue of osmotic pressures is important in biology. And as I told you several times, I'm no expert in biology. So I don't want to talk about it because it's too risky to say wrong things. But I can point you to a nice YouTube lecture where you can learn about osmosis in biology. And it tells you how cells receive-- exchange across the membrane molecules, depending on whether there is a chemical potential.

If there is equality of chemical potential, of course, you're in equilibrium. But if the concentrations are different than those that you are at equilibrium, then depending on the sign of the difference in the chemical potential, you can have a flow in one direction or the other. And that is helpful in regulating the metabolism of the cell.

And in this video, you also are brought to understand what the difference between plants-- membranes of plant cells and animal cells are. They are kind of different. And they behave in different ways. In the plants, they can sustain more pressure. In the animal cells, they are more fragile, so to speak. And they can break.

There are also important applications in chemical engineering. For example, recently, of course, there is this big topic of developing lithium-ion batteries. And here, you have these lithium-ion molecules, these pinkish molecule that are in dilute solutions of solvents. And these solvents, for example, in this case are some kinds of carbonates.

So the osmotic pressure-- of course, if you are at the osmotic pressure, you are in equilibrium and nothing moves, on average. But if you do not have the concentrations that deliver the osmotic pressure, then you can have osmosis, which is the movement of the ions in search to a better equilibrium.

So let's talk about seawater as the example. Seawater, on average, per kilogram of seawater has 35 grams of salt. And of course, there is this fact that the salt dissociates. And so you get ions of sodium and ions of chlorine, chlorine that gets hydrated, meaning that the molecules of water which are polar get attracted from one side.

The oxygen side is negative. So it's attracted by the sodium, which is a positive charge. And the hydrogen side is positive and is attracted by the chlorine ion. And so you get these molecules. Typically, around six molecules, on average, are around each ion.

However, not all ions get so far apart. Sometimes, they get also-- they form ion pairs. In other words, they also attract-- still attract each other. And they are still also surrounded by water molecules.

So suppose we apply the van 't Hoff equation by neglecting these ion pairs and assuming that 100% of the salt goes into these hydrated ions forms. So that means that now I have to do the counting in order to see how much solutes I have.

So for each molecule of salt that I put in solution, I get two molecules that are the ions. But in some sense, you should also account for the fact that these hydrated or solvated groups-- since they like to stay near each other, they could be also thought as a big molecule. So part of the water molecules that are in solution and get engaged in these solvation groups-- and so they're not available anymore as to play the role of solvent. So the mole fraction of water also changes because you lose six molecules per ion that is formed here.

So let's see. So the number of molecules of water initially are the mass, which is 1 kilogram, minus the 35 grams. So it's 965 grams. And if you divide it by the molecular weight, you get how many moles of water you have initially. And then for the ions, their amounts are the 35 grams divided by the molecular weight of the salt, of the salt, not the ions, which is-- gives this much more. So these are how many molecules of salt you have put in. And then you assume that the mole fractions-- that as they go in, you double this amount.

So in the end, you can compute the mole fraction of water and-- minus its logarithm. You can check, of course, that the Taylor series approximation is good because this is a small number. The mole fraction is 0.978. So it's very close to 1. So that approximation-- but if you do the calculation, there is a little difference. $1 - y$ is a little different from $-\ln y$.

So let me just keep the $-\ln y$ without that approximation, which-- since we can do logarithms. And we can compute, therefore, the partial pressure-- I mean, the osmotic pressure using van 't Hoff.

RT -- I assume 298, just to put a number here. The specific volume-- here, it is correct. This is with a double v . Double v is the specific volume of pure water. So it's 10 cubics meter, 10^{-3} cubic meters per kilogram. But since we need it in moles, we have to multiply times the molecular weight of water.

And in the end, you can compute this in terms of kilopascals or bars, depending on what makes you have a better feeling of how much this pressure is. Atmospheric pressure is 1 bar, more or less. And so the osmotic pressure-- this osmotic pressure is 30 bars. You can measure it in terms of a column of water. The weight of the column of water gives a pressure. It's 300 meters.

Now, it turns out that this number is not precise because there-- we've made some approximations here. For example, we have neglected these ion pairs. The true number-- first of all, it varies from place to place. For example, in the Mediterranean, it's higher than in the Atlantic or in the Pacific or in the Arctic-- is even lower.

So there is a range of values of salinity that go between 30 and 38. And the measured osmotic pressures are between 220 and 290 meters. Still, it's not too bad. But we can do better.

So here is-- sorry for the fine print. I will put on Canvas this Excel file in which I've done the calculation. So this first column gives that calculation that we have just seen. And that leads to 310 meters of osmotic pressure.

And we have assumed 100% hydration and 0% ion pairs. And also, we have assumed no water. I forgot to say it. But we assume that I'm not counting for the water molecules that are lost because they go and formate these groups.

Now, in this other column, I do the same calculation by assuming six molecules per ion, and also by assuming that some of these, 18% of the ions, are grouped in ion pairs, which are also surrounded by six waters. If you take that-- why did I take 18%? That was just my only fitting parameter in order to get the measured value of 288. So if that were the only effect, then my model would get the actual measured value by assuming this percentage of ion pairs with respect to the hydrated ions.

And these other two columns give you the same calculation with the same parameters. But in the Mediterranean-- for example, in Naples-- you have this much, 316, whereas in the Arctic Ocean, it's 242 because the-- because of the-- changes the percentage of salt that is dissolved there.

Now, first of all, let's view it in this way. So imagine you have a long pipe. So this is air. This is z equals 0. And this is z equals minus h .

So if the density of seawater is approximately equal to the density of water and if I have here the pressure P_a , here, what pressure do I have? That's the static Bernoulli equation or whatever you want to call it.

So now, this is a porous plug semipermeable to just water. So suppose I want to have this level. This is called-- let me call it minus h_1 . And this is h_0 .

So on the outside-- so below here-- I have the pressure that I just said. In the seawater at that level, that's the pressure. And it's h_1 . On the other side of the membrane, I have a pressure that is given by-- I assume that also here, the pressure is equal to atmospheric. It's a bit higher because of the column of air. But we can neglect the density of air, of course. And so here, it's P_a plus ρ . It's, again, assumed that it's the same ρ , gh_0 .

So across the membrane, I have a pressure difference. So the pressure difference is ρ g h_1 minus h_0 . And if I am at equilibrium at the temperature that prevails there, that is equal to the osmotic pressure. So that would be the equivalent of ρ g times 288 meters.

So there is a huge difference in pressure there. So this porous plug must be able to sustain that pressure. And in fact, this is why. As you know, in many countries, the production of freshwater from seawater or from other saline waters is done precisely this way. It's not really done in this-- you can do it in a more practical way, this, by having-- you take some water. You take a piston. And then you push. You generate a high pressure here. So this is seawater. And out on the other side of the membrane, you get freshwater.

In order to-- for this to happen, you have to impress a pressure. If you want freshwater at atmospheric pressure here, you need at least 288. You need more. If you just do 288, you get a drop. If you put higher pressure, you get a flow.

Of course, you need a membrane that is capable, like in the new graph-- capable of sustaining those pressures. And that's why you need-- part of the job is to make the membrane. And part of the job is to sustain the pressure so that the membrane can resist that pressure. And that is how you make seawater.

Now, you could also ask yourself, what is the minimum work needed to do this job, because this is a very energy-consuming activity? And there are various-- there are more than 20 different methods that people have devised in order to separate freshwater from seawater because, of course, it's a matter of survival. People have developed many ways, even in the old times, when you didn't have electricity.

Each of these methods has an efficiency. And if we are talking about second-law efficiency, like we like to do in this course, you have to compare what you spend with the minimum you could have spent. And so what is the minimum that we could-- we need to spend in order to produce 1 liter of-- or 1 kilogram of freshwater?

And, well, you have it in front here because suppose I had this long pipe. And this is my boat. And the guy takes a bottle of-- an empty bottle, pulls it down-- or a bucket. Maybe it's easier. Take the bucket. Fill it with this water, which is fresh because it's on the other side of the membrane, and pull it up. And then after a while, this will be replenished. And you can get some more.

So how much work do you need to do that? It's just the work for lifting from here to there. So it's ρgh , essentially. And it turns out that it is about-- I think I have it in the-- I think it's about 1-- or I mean-- well, g is about 10. ρ is about 1. So it's about 20. And this, I think, is kilojoules per kilogram. Is it kilojoules or is it joules? That's the question.

Because ρ is, say, 1 kilogram per cubic-- I'm sorry. ρ is 10 to the 3 kilograms per cubic meter. g -- for simplicity, let's call it 10 meters per second squared. And then the other is 288 meters. Anybody wants to help?

[CHUCKLING]

So this is kilogram per-- that's newton-- kilopascals. No, this is pascals times 10 to the 3-- is kilopascals per meter. I think it's kilojoules. Well, you check it.

There are other ways in which, for example, if you didn't want to put-- because this is a commercial product. But if you go and look up a solar still, you find an inflatable survival object, which has something like this. So this is an inflatable donut-shaped thing with a plastic dome. And here is the seawater.

So you get sunlight coming in, which makes some of the water inside evaporate, like in a greenhouse. The water vapor that evaporates-- the salt remains here. So what evaporates is just water. So the water vapor goes and forms and condenses on here because you also, typically, have some breeze or some wind that will cool that dome. And then the drops by capillarity stay and get collected on this side here.

So eventually, you get some water on this side. And you put a little pipe and a little bottle, and you collect your water. So that's a survival kit for the fishermen. And what kind of efficiency would that have? That will also help me figure out if this k should be there or not.

How much energy, how much solar energy, do I need in order to evaporate 1 kilogram of water? It's the enthalpy of vaporization of water, which is of the order of 2,200 kilojoules per kilogram. So I take this k away because I know that the answer is that that thing has an efficiency of 0.1%. So it consumes 1,000 times the minimum.

These ones, the commercial-- let's see. Do I have it somewhere? Yeah. The commercial ones-- these are the most typical, most important industrial technologies. One is called pressure-retarded osmosis, which is exactly like that. And the other is called reverse electrodialysis, which uses the fact that since the ions are electrically-- so if you put an electric field, you can separate the ions and, therefore, take them away from the water. So that's one way of separating, which is less efficient.

There is a factor of 2 of inefficiency there because you have to separate the two ions for each amount of water. And so I think that is the reason why there is this factor of 2 in the theoretical maximum efficiency that-- or minimum power that you need to spend. Here, it's for-- to treat 1,000 cubic meters per second.

The reason I was interested in this is because of the fact that you see here, you have freshwater in equilibrium with-- so this is an equilibrium situation. So that's the freshwater. If you add some freshwater from above, that freshwater-- this will add to the pressure on the membrane. And therefore, the water will flow into the ocean. So that's one way. That is-- would be an intelligent way to put freshwater in the sea.

So when do we put freshwater in the sea? Why would you want to do it. Not the fisherman, because he wants to drink it. But the rivers-- rivers bring freshwater to the sea. And usually, they just dump it at the same level that's dissipative.

The best way would be to take that water, for example, one bucket at a time. I take that water from the river. I put it in my bucket. Then I let the bucket go down. So while doing that with a pulley, I have an engine-- will produce electricity. And that is-- we know how much work. We get this much work per kilogram of water that we put down.

And then at this level, when it is at this level, it will cross the membrane spontaneously. And I don't need a big difference in pressure to make it flow. Of course, I do need some difference in pressure. Otherwise, the flow will not be able to win the pressure drops in the porous plug. So I cannot obtain 100% of this. This would be the zero-power situation. If I need power-- so I need to do it at a certain rate-- I need, instead of being here, let's say, being halfway.

So I use this much head in order to push the freshwater into the sea. In any case, this is what is called blue energy. And it is an enormous potential. And it is, of course, a renewable kind of energy.

It has an enormous potential if you look at what the rates are for rivers. So even if I take just this small contribution for the-- for Po River in Italy, which is only 6,000 cubic meters per second, you see with a pressure-retarded osmosis power plant, we could make 6 times 3,000 megawatts. So that's 18,000 megawatts. That's a lot.

And people have tried to do this. There are pilot plants. But the problem with these technologies that have to do with-- that need membranes is that membranes get clogged. Freshwater from the river is not really just water. There is mud. There is dust. So you need to keep these membranes clean. You need huge surfaces in order to treat these kind of flow rates. And therefore, up to now, it hasn't developed yet in any reasonable even small fraction of the electricity produced.

This map also shows the salinity differences in the various oceans and seas. So here is the Mediterranean. That's about 38. Here is the Arctic-- is below 33.

Now, the idea of-- this idea-- using the pressure-retarded osmosis is this. There was a power plant in Tofte, Norway that tried to do it. But I think they closed. I think that there is somebody else who is starting another project in trying these things.

So you get water from the river. And instead of putting it directly-- so I put a nice wall here. And the water does have to go a different path. It has to go into these pipes. So these are pipes made with a porous material, permeable only to freshwater, to just the water molecule.

So what I do is to put inside-- on the inside, I put seawater, pressurized, pressurized, not to the osmotic pressure because otherwise, we would have no flow, but to a pressure which is halfway the osmotic value. That way, there will be the driving force that would push water through the porous medium.

And so this pipe will need to be in-- with an enlarging cross-section so as to accommodate this increasing flow of water. So here is the seawater that gets in. And what gets out is the seawater plus the water, freshwater, that has been added. And then it goes here.

Now, of course, in order to pressurize the seawater, I have to consume more energy with a pump. And that pump is-- the work-- the power is the product of the pressurization that I have managed to obtain times the volume flow rate. And there should be a dot there.

What flows out is, again, at that pressure, perhaps a little less because of pressure drops. And so I don't just put it in water. I put it in a turbine, hydraulic turbine, which is going to elaborate, though, a higher volume flow rate because it's not only recuperating the volume flow rate that I have pumped, but it also gets the volume of freshwater that is being added.

And so it's on the difference between this volume flow rate that I gain. And again, the net power is going to be the difference between what I obtain from the turbine and what I spent from the pump. So in principle, this is relatively simple. But again, the problem is these porous medium and porous membranes.

This number that we obtained here is the work of separation of some solvent from this dilute solution. In general, the problem of separating stuff is important for many applications. For example, in the nuclear power industry, you have to separate-- in order to concentrate uranium-235, you do-- you need to-- you make it into gas-- hexafluoride of uranium. And then you separate the gas using the fact that the molecules are different because they have a different-- they're different isotopes, the 238 from 235. And so that is consuming energy in order to do this separation.

So we know that we did something very similar earlier. So I'm going to go quickly on these energy balances and entropy balances. But the idea here is that we have a mixture, in general, given composition. We do it in a given temperature and pressure. So now I want to operate at a given temperature and pressure, which is-- could be, for example, the atmospheric reservoir. And I want to obtain fully separated, pure components.

Of course, in order to do that, I need to spend some work that's represented by the weight that goes down. So I do volume balance for the overall system, including the reservoir. Here I call B and A. B is the final state. And A is the initial state.

So the energy balance, the entropy balance, and then the equation for the stable equilibria of this reservoir with variable volume, which is this-- then I combine, as usual, these equations by eliminating the reservoir variables. And here, I'm also using the fact that the final-- since we are in equilibrium with the reservoir, it fixes also the final temperatures of the pure components at the end. So T_B , T_A is equal to T_R .

And the linear combination of these balances yields this expression for the work, which is one of those characteristic functions. I think it's the one that we call with the letter ϕ .

And then there is a plus T times the entropy generated. So if I do it in an irreversible process, this will make me consume more. Well, let's talk about the limit in which I do the best. So this is not there. So that would be the reversible case.

And then I recognize that this linear combination of properties is just the Gibbs free energy. So essentially, the minimum work for separating at constant T and P is the difference between the Gibbs free energy of the final separated state and the Gibbs free energy of the initial mixture.

And now, since we know how to do things, we know how to-- we know, first of all, that the Gibbs free energy-- now, in general, it's not an additive property. But if all the systems are at the same T and P , then the Gibbs free energy is additive. And so in the final state, the Gibbs free energy of the composite is equal to simply the Gibbs free energy of the sum of the Gibbs free energy of the parts.

And again, here I've used the fact that the Gibbs free energy of the pure component is also equal to the chemical potential of the pure component, whereas for the mixture, the Gibbs free energy due to the Euler relation is related to the chemical potentials. So it's the weighted sum of the chemical potentials.

And up to here, we made no assumptions. If we further make an assumption of either being in the case of a Gibbs-Dalton mixture of ideal gases or an ideal solution-- and this time, it's ideal for all components, not just the solvent-- then if we apply the rules that we have seen so many times now, this gives us a formula, which is similar to the entropy of mixing.

So suppose, for example, that we want to separate air. We want to separate air into its components. Now, air is a mixture of many gases. It's not just nitrogen and oxygen. The typical composition at sea level is this one-- 78% of nitrogen, almost 21% of oxygen, 0.9%, almost 1%, of argon, and 0.04%-- these are the 400 PPMs of CO_2 .

I want to compute this-- the summation. And it's simple. And then I want to divide it also by nM so that I get the result per kilogram of air. So I have to divide by the molecular weight. And the molecular weight is the weighted sum of the molecular weights of the various constituents weighted on the mole fractions. So it's 28.97 with this composition. And then I choose a temperature-- say, 298 K. So the minimum work to separate air into its components, 1 kilogram, is 48.5 kilojoules.

We could have modeled air in a simpler way. Suppose air is just made of-- suppose we have a mixture in which I have only nitrogen and oxygen with the same relative composition as they have in the actual air, but there are no other components. If it were only those two, then I do the same calculation. I see that the work for separation is 44.

It's quite a bit of difference. And these additional about 4 kilojoules per kilogram here are needed because here, I have insisted in separating all, including argon and CO₂. And those are very few molecules. And they require a lot for separating per unit of amount that you want to separate.

So this formula applies not only in the gas phase. You can also apply for an ideal solution liquid phase. For example, let's take a mixture of water and ethanol, which, at the appropriate mole fractions, is like wine-- a model for wine, of course. It's not just these elements.

So if I take a wine 14 degrees by volume, that's-- ϕ is the mole fraction-- it's-- is the volume fraction of alcohol in wine. If you do the calculations to go-- to pass from volume fraction to mole fraction, it's 4.7 mole fraction. And if you want to separate entirely the water from the alcohol, you apply this formula. And again, if you want to do it for 1 kilogram of wine, it's 24 kilojoules, what you spend.

Now, the question is, well, if I want to separate-- for example, if I want to produce oxygen because I need it for medical applications or I want to produce nitrogen because I need it for agricultural applications or if I want to produce argon, which is needed in the lab, or if I want to separate CO₂ from the atmosphere because I want to segregate it, which is something that people are considering these days, how much do I do? And do I have to separate entirely my mixture?

Well, the answer is no. What you could do-- you could do exactly just what you need, not-- or if you do something-- if you separate something that you don't need, then you'll spend energy to do the separation. And you're left with something that-- since you don't need it, you mix it again. So it's wasteful.

So concentrate on what you need. So for example, take your initial mixture-- say, air. And suppose you want to extract some oxygen. So the best way-- the minimum expenditure would be to extract a little bit of oxygen, not the entire amount of oxygen in your air.

So you treat the volume of air and extract only a fraction of it, a fraction of the oxygen that is present in that air. So you put out air, which is leaner in oxygen. And you have your oxygen here. So essentially, this is what we're doing. We take lots of oxygen, little oxygen. And then the little you separate. And the lot-- you leave it with the air. And the other components-- you leave it there.

You do the energy balances and the entropy balances and the volume balances again. And again, you find that the difference-- I mean, the minimum work of separation is given by the difference in the Gibbs free energy of the final situation minus that of the initial mixture.

Now, the final situation or the final composite of mixture plus the pure has, therefore, two contributions. This is for the mixture, which now has a final composition, y mole fractions B, which are different from the mole fractions that we started with, plus you have the Gibbs free energy of the pure component here. And again, remember that chemical potential is, for the pure substance, the Gibbs free energy. And again, we are using additivity for Gibbs free energy because the two subsystems are at the same temperature and pressure.

For the mixture here, I would have to write that the-- from the Euler relation that the Gibbs free energy is simply the summation of the amounts times the chemical potentials. I have to be careful, though, that when I get to component i in the summation, there is this extra little piece.

So I'm going to write this summation in two pieces. One is with the n 's that remain unchanged. And the one that is separated, its contribution is multiplied by the same chemical potential. But I put it outside. So component i -- I will find it outside of the summation for the amount that I separate and inside the summation for the amounts that remains.

Then I use ideal solution or ideal Gibbs-Dalton mixture behavior so that I can express the chemical potentials of the mixture in terms of those two component times $RT \log$ of y 's. And doing that substitution, I obtain this expression. So this is the general expression for the work in terms of the initial amounts and of the amount that I have extracted of component i .

Now, there are several cases that may be interested-- interesting. For example, I could consider the complete removal of constituent i -- so a case in which I do want to take, for example, CO_2 . Maybe I want to extract it all. So I don't want to be that-- to be left with any in my final mixture. So n_i should be 0. That's complete removal. And that expression here gives us this result.

And there is an interesting possible limit, which is one in which the mole-- initial mole fraction of my component, like, for example, the CO_2 case-- the initial mole fraction is very small because CO_2 , we said, is 0.04%. So that means very less than 1.

In this case, the second contribution goes away. And so the minimum work is just minus log of the mole fraction, the minimum work per unit amount extracted-- or you could use a partial removal in which I don't extract it all. I extract only a fraction. I could call this fraction χ . The amount that I have initially-- I want to-- say I want to extract only a fraction χ of the initially present amount.

So that expression rewrites this way. And here also, there is a limit, which is the same result, but which obtains in the opposite situation, when I extract-- when I separate only a minute fraction of the originally present amount. This is typically used in the oxygen or nitrogen separation case.

So if we do, again, assume the same air that we considered, standard conditions-- and suppose we want to do a minute or minor removal so that this fraction χ is small. So I can apply this formula here.

So if my objective is to extract 1 kilogram and these are per kilogram of the molecule that I want to extract-- so I want to extract 1 kilogram of nitrogen. I spend 21.9 kilojoules. If I want to extract 1 kilogram of oxygen, I spend 121. It's more because oxygen is more rare-- is by 4 times less. If I want to extract argon, it's 290. If I want to extract CO_2 -- is 440.

Of course-- I want to make sure that I cannot build a perpetual motion machine of the second kind. That would be something in which I use this method of separation. And then once I have separated, then I mix them again. And we know that by mixing, you can extract work.

So clearly, if I start from air, I separate it. And I spend so much to separate it. Then when I mix it again, the maximum work I can separate must be equal-- I mean, I can obtain must be equal to the minimum work of separation.

So for example, if I take the right amount of nitrogen that is needed to form again the air in its original composition, I have-- this is the right-- that right amount. So I have to multiply these ones times those amounts. And then, of course, the sum of these works properly weighted by those amounts gives me the same work that I can obtain by mixing, which is-- so this 48.5 is the same work of separation that we have used in the complete separation case in the previous slide. So this is just a consistency check.

Also here, I have another Excel file that I will share in which these calculations are done. This is where I took those results. And for example-- this is the partial removal, partial removal separation. These are the amounts of the mass percentages of the various constituents in air. So if I take the weighted sum of these times that, I obtain the same as if I did the complete separation.

Now, this also is in the same or another Excel file. There is this calculation of the alcohol content and the separation for wine. It's interesting that-- takes a while. You may have to be careful with that. Volume fractions, mole fractions, mass fractions can get confusing. But eventually, you figure it out.

So this is already the result that I mentioned before. What was the number? 21 kilojoules per liter if I take wine. This is 12 degrees, 12 volume fraction of ethanol, 12%.

This other thing is an application of the partial removal case. Sometimes, winemakers in the north of Europe, where there isn't so much sunshine, the alcohol content, the natural alcohol content, that you obtain is low. And you have to find ways to raise it.

One way is to buy wine. It used to be that they-- it used to be cheaper to buy wine from the south of Italy or from Greece or from Spain because there, you get a lot of sunshine. And so the alcohol content is high. And it used to be called the cut wine, where you need to cut the poorer wines of the north. But now, also the people in the south of Italy, Greece, and so on have learned to make good wines. And so they don't sell it for low anymore.

So an alternative is to separate, take a little bit of water away from the-- so for example, suppose you have a 12-degrees wine and you want to bring it to 14. All you need to do is to take some water away.

Of course, you have-- you cannot boil it. Otherwise, you're going to ruin the wine. But maybe you can flush it-- low pressures so that at room temperature, you get some-- so there are ways in which you can extract without ruining the rest of the good part of the wine. And so this calculation gives you-- with the formulas that we've given, it-- all you need to do is to figure out, how much water do I have to take away?

And so this is the Δn_1 that I need. If you look inside, there is a formula that shows you how I calculated that. And once I determine that, I can apply-- I can find the χ -- so the cut fraction that we defined in the slide. And we apply those formulas. And it comes out that we need to spend 0.86 kilojoules per liter in order to make this little improvement in the alcohol content.

And of course, therefore, you'll find that the wine-- if you do this way, you find that the winemaking becomes also an energy-- how do you say-- an energy-consuming industry. Yes.

So let me move on to another interesting application. It has still to do with air. It has to do with stratification of air. As you know air is stratified. As you go up to high mountains-- for example, top of Mount Everest-- you can read books about what they call thin air.

So why is it thin? It's thin for two reasons. It's thin because the overall pressure is low, maybe one-third of the atmospheric pressure at sea level. And it's thinner also because the oxygen content percentage-wise is less. Why less? Because the oxygen molecules are molecular weight 32. Nitrogen molecules-- 28. So the oxygen is heavier. It likes to stay more towards the sea.

So the stratification is due to these differences in molecular weights of the molecules. And since we are so interested in CO₂, CO₂ is 44 molecular weight. So even more, there is a stratification.

For example, when we say, ah, the CO₂ concentration at Mauna Loa, which is the place where they measure CO₂, it is a mountain. That mountain is at 3,400 meters altitude. It's also maybe representative of the content of - the average content in the atmosphere.

But the point is that if there, there is 426 PPM, for example, today, it means that at sea level, you have 500 and something. We'll calculate it in a moment. There is a big difference due to these differences in molecular weights of the molecules.

So how do we go about and do this calculation? We need a model, as usual. And the model that we make is very similar to the model that we have made for bulk flow.

You remember, for bulk flow, we said if you have a fluid element which finds itself at elevation z in the gravity field and which has a velocity-- let me call it w -- and it has a mass m , when we do fluid mechanics, this is a fluid element or a fluid parcel. And we assign the properties in the local equilibrium assumption. We assign the properties of this fluid element by assuming that its energy is the energy that it would have with the same composition, volume, and entropy in a stable equilibrium state, plus we add these mechanical effects that are the kinetic energy and the potential energy.

The next time, we'll assume also that this particle has a charge, electrical charge, and that finds itself in, for example, an electric field. An electric field is really minus the gradient of an electric potential. And so here, we will add q times ϕ . It's mechanical kind of effect. So the thermal aspect of the properties are still contained in the fundamental relation for stable equilibrium states. And these are other additional contributions.

Now, for the problem of stratification in the atmosphere here, all we need to do is to forget about motion. So we are talking about the still atmosphere-- of course, it's a model-- and not charge.

So we have this potential energy. And m is the mass of the fluid element, which may vary. We're going to assume that it may vary here because since the composition changes with elevation, also the mass changes.

So the mass is given by the moles times the molecular weights. So this is sea level. This is the generic layer at elevation z , where we assume that there is a temperature, a pressure, and the composition. And there is a adjacent layer, which-- at z plus dz , with a different temperature, different pressure, and different composition.

Now, when we have this kind of energy, you remember-- and when talk about-- when we derive the conditions for mutual equilibrium, the conditions for mutual equilibrium were equalities between some partial derivatives of the energy of the fundamental relation that, precisely because they appear in these conditions for mutual equilibrium, are important. And that's why we give them names-- temperature, pressure, chemical potentials.

But here, when we take the derivative of energy with respect to S and V , we still get temperature and pressure. But in a model like this, when we take the derivative of E with respect to the amounts, I have the amounts here, which give me the chemical potential contribution. But I have the amounts also inside the mass, which gives me an additional contribution, which is a gravitational contribution. If I have the charge, which we'll do next time, we'll have also an electrical contribution.

So this more general kind of potential that so far we've called chemical-- we will call it the total potential because it has several contributions. And when it is only chemical plus electrical, it's the electrochemical potential.

Today, for the stratification problem, you could call it-- what is it, electrochemical-- gravitochemical potential. Whatever. The name doesn't matter. What's important is that this total potential is given by the chemical potential. So it's the derivative of this one plus the derivative of m . And so that brings up the molecular weight of the molecule with respect to which you are taking the derivative.

So I think we are done today. But the idea is that we will obtain our results and, therefore, the distribution in concentrations by imposing chemical potential equality between layers because the layers are attached to one another. So they are in mutual equilibrium.

In practice, they are not because as you know, the temperature in the atmosphere is not uniform. We will assume that this change in temperature in the atmosphere can be neglected locally, but not globally. So that will be a little bit of an assumption.

So we assume that we make an equilibrium between these two, assuming that T at z plus dz is equal to T at z . But we will not assume that the chemical potentials are. And that's how we will get our results.

But the time is up. And so I think I have to simply wish you a good spring break and see you in 10 days.