

Lesson 10: Liquid-gas phase transition

Notes from Prof. Susskind video lectures publicly available on YouTube

Introduction

We all know about phase transition from ice to water, then from water to steam. It has been understood qualitatively for more than a hundred years. Van der Waals¹ explained it from Boltzmann statistical mechanics.

We are going to see how phase transition from liquid to gas works. We will not follow the way Van der Waals approached it. His approach was very elegant, very nice, and equivalent to what we are going to do.

Our approach, however, will be based on a system that we have already studied: magnets, more specifically the Ising magnet. We will see in what way the Ising magnet² is a model for the liquid gas phase transition.

Before we do that, let us review quickly what we learned in the last chapter. We studied the mean field approximation, not for the liquid-gas but for the Ising model. It enabled us to see how and why phase transition can occur in certain systems. We also know that it does not work in one dimension. In fact it then gives the wrong answer. Indeed, before studying several dimensions, we showed that in one dimension there was no phase transition.

¹Johannes van der Waals (1837 - 1923), Dutch theoretical physicist and thermodynamicist.

²Remember that the Ising magnet was more a simple mathematical model for any kind of phase transition, when elementary building blocks have two possible states, and adjacent pairs are coupled in some way, than a real model of magnetism in nature. It is too crude for that. But it is quite adapted to illustrating how sudden changes in the properties of certain systems can appear when some parameters are changed continuously.

The mean field approximation always tells us that there is a phase transition. We saw why it is appropriate to use it when the number d of dimensions is large. We also stated, and gave some heuristic justifications, that there is a phase transition as soon as $d \geq 2$.

By the time we get to three dimensions, every particle in a lattice has enough nearest neighbors – six of them in the basic cubic lattice case – that the mean field approximation becomes very good. It is not accurate in the sense of 1% but accurate in the sense of probably 10%.

So the mean field approximation is a good way qualitatively to understand how phase transition takes place.

Let's consider a two-dimensional square lattice. At each site, also called vertex, there is a little magnet that can be up or down, figure 1.




Figure 1: 2D square lattice of little magnets, sometimes called spins. They are located at the sites of the lattice. Each spin is either up ($\sigma = +1$) or down ($\sigma = -1$). And the edges between the sites are also called the bonds, or the links.

Each little magnet either points up, in which case its cor-

responding σ is equal to $+1$, or points down, in which case its corresponding σ is equal to -1 . The magnets are positioned at the sites, and they are linked by bonds, also called edges, or links.

For each magnet, σ is a random variable. For the i -th magnet the random variable is more precisely σ_i . They all have the same distribution of probability, and therefore in particular all the same mean $\bar{\sigma}$. When $\bar{\sigma} \neq 0$, by definition there is magnetization.

We talked about two different kinds of models for the magnets. Before the Ising model, we first considered a very simple model where the little magnets have no relation among each other. We say that they are uncoupled.

They can be pictured as forming a big linear sample, or a multidimensional one. It doesn't make any difference in this model. Moreover the whole system was plunged in a magnetic field B . Then there is no magnetization only when the temperature is infinite.

We saw that there is of course magnetization at $T = 0$, because then all the elementary magnets must point in the same direction, which, taking B into account, makes the energy of each of them minimal. As soon as the temperature is less than infinite there is some magnetization. And there is no phase transition.

Then we turned to the Ising model, which is more difficult. We studied it first in one dimension, then in several dimensions. In all Ising models, anyway, the energy in the system is no longer stored independently in each spin, but is stored

in pairs of adjacent spins. Equivalently we can say that it is stored in each bond, or link.

The energy in a bond is lower if the two spins are parallel, and higher if they are antiparallel³. If two neighboring spins are *parallel*, we call that an *unbroken bond*. Then the energy of that bond is low. If they are *antiparallel*, we call that a *broken bond*. Then the energy of the bond is high.

We wrote an expression to represent that. The energy of the whole system is

$$E = -j \sum_{links} \sigma_n \sigma_m \quad (1)$$

We don't try here to make the notation fancy⁴. The values

³We saw that we could choose the other way around. It is somewhat more intricate to follow, but is absolutely equivalent.

⁴Remember that in physics, as well as in mathematics, notations are only shorthands for descriptions and explanations in plain English. What makes them sometimes a bit cryptic is that, whenever we can, we devise them in such a manner that they lend themselves to algebraic manipulations and rules which, more or less automatically, express operations, like if $x + y = 8$ then $y = 8 - x$.

When he was ten years old, the second author one day witnessed academic inspectors come to his math class and ask: if $x = 5$ and $x + y = 8$ what was y ? Since all the pupils who were not already bored said y must be 3, the inspectors concluded gravely that we were ripe for higher algebra. This unfortunately stressed the mechanical rule to the expense of common sense, and contributed to the esoteric aspect of mathematics, which many people etched forever in their minds as their principal feature.

All too often official education consists in teaching a formalism which makes obscure what is simple. That is also true in physics. And that is why, in these courses, we explain higher physics as plainly as possible, with models described as straightforwardly as we can, and the least amount of arcane equations.

σ_n and σ_m simply mean the two ends of one link. And the sum of over all the links, in figure 1, or whatever lattice we are talking about, including the one-dimensional one. The constant j is just associated to how tightly the spins are bound, how hard it is to break a bond.

How much energy does it cost to break a bond? If we start with two aligned spins, then $\sigma_n\sigma_m$ is $+1$. If they are antialigned, the product is -1 . So the difference, to go from $+1$ to -1 , in the sum in expression (1) is two units, with a negative sign. But there is a $-j$ in front. That means every broken bond costs an energy $2j$.

Then we introduced a possible term in the energy that would correspond to a uniform magnetic field⁵. It is a term which acts not on neighbors, not on links, but on the individual sites. It is an extra energy if a site is up or down. It doesn't depend on what its neighbors are doing.

So let's add it in equation (1). The total energy becomes the sum of two terms

$$E = -j \sum_{links} \sigma_n \sigma_m - \sum_{sites} h \sigma_n \quad (2)$$

In the second term, if h is positive, then $-h\sigma_n$ is lower when σ_n is positive.

In the calculation of E , if the second sum is all there was – which incidentally is the first model we studied –, then

⁵In other texts, the uniform magnetic field is often denoted H or B . And we have already used both notations in the past chapters. But in this final chapter, for the nice symmetry it makes with j , we are going to denote it h .

it would favor each σ_n being positive, or being up. That would lower the energy of each spin σ_n . And of course, with h negative, or changing the second sign in equation (2), it would be the other way around.

Nevertheless thermal fluctuations will cause σ_n to randomly point down sometimes, And at high temperature, it will be down as much as up.

So, in the complete expression (2) for the energy E , the first term wants the spins to point in the same direction, and the second term wants them to point up.

Notice that, since it is only the difference between various levels of energy which matter, we could add a constant k on the right-hand side of equation (2) without changing anything. You may also observe that the higher j is, the harder it is to break a bond. And the higher h is the more energy is stored in one spin when it is down. But let's not pay attention to that for the moment. We are in a simple abstract model to study phase transition, and j and h are just given numbers.

How much energy do we have in the first term if all the spins are aligned? That is the lowest possible energy coming from the bonds. It is just the number of links times $-j$.

Incidentally, how many links are there in a 3D cubic lattice? Every site has 6 neighbors⁶. But, if we want to avoid double counting, it is the number of sites multiplied not by 6 but by 3.

⁶We neglect the border effects if the size of the whole lattice is big compared to the size of the bonds.

It is easy to see in 2D, figure 2. Starting from the lower left site, say, σ_1 , draw two links. That leads us to the two sites σ_2 and σ_3 . From each of these sites, there are only two new links, etc. So altogether, without double counting, each site corresponds to only two possible new bonds.




Figure 2: In a 2D square lattice, count of the number of bonds compared to the number of sites.

The same is true in 3D, with three new bonds for each site. And generally speaking, the number of bonds is the number of sites times de number of dimensions.

In any case, looking at the two terms on the right-hand side of equation (2), we are in the ground state, i.e. the state with lowest energy, when everybody is aligned, and in the direction of h . That energy is a certain quantity which we could subtract from the right-hand side of equation (2), to make the ground state energy equal to zero if we liked.

However, as said, the actual energy of the ground state is unimportant. Only differences in energy levels are important in statistical mechanics. Here what is important is the difference in energy between a spin up or down, and an un-

broken or a broken bond, and how they sum up.

Let's remember what we found out in the mean field theory. We calculated the partition function *for one spin*. It has the usual form

$$Z = \sum e^{-\beta E} \quad (3)$$

where the sum is over the possible energy levels attached to that spin. Each energy level is a part of the right-hand side of equation (2), which was for the whole system. For one spin, an energy level is

$$E = -j\sigma \sum_{links} \sigma_m - h\sigma \quad (4)$$

Then, in order to make equation (4) tractable, we used the *mean field approximation*. Let's go over it quickly, illustrated with a picture in 2D, although in truth the number of dimensions should be large. So we consider one spin, or site, at the center of figure 3.




Figure 3: One site and its neighbors, leading to the mean field approximation.

The center site is surrounded by a bunch of neighbors, four

in 2D, six in 3D, and 2 times d in d dimensions. The spins of the neighbors have an average value. That average might be zero, or it might be a net non zero value. The more neighboring sites we have, that is the higher the number of dimensions, the better approximation it is to just say that the average of the neighbors is $\bar{\sigma}$, which is the average of all the spins in the system.

Moreover we can say, without making a statistical mistake, that the spin under consideration sees a collection of neighbors, each of which has spin $\bar{\sigma}$.

So the energy of all the bonds in figure 3, is

$$-2dj\bar{\sigma}\sigma \tag{5}$$

where d is the number of dimensions, j is the usual factor measuring how difficult it is to break a bond, $\bar{\sigma}$ is the average spin in the system, and σ is the spin under consideration at the center of figure 3.

Secondly, what about the sum over sites on the right-hand side of equation (2)? Well, that term doesn't care about neighbors. For one site, it is just the extra $-h\sigma$ appearing in equation (4).

So, for one site, the complete energy level E , as a function of σ , is

$$E = -2dj\bar{\sigma}\sigma - h\sigma \tag{6}$$

And, with the average spin $\bar{\sigma}$ being just a number, we pretended that this is the statistical mechanical problem of *one*

spin, the energy of which is given by expression (6).

Of course, equation (6) is simply

$$E = [-2dj\bar{\sigma} - h] \sigma \quad (7)$$

That is the calculation we did when we first looked at these magnets. And we found out that *the average of the spin under consideration* in the center of figure 3, which we called at first $\bar{\bar{\sigma}}$, satisfies⁷

$$\bar{\bar{\sigma}} = \tanh ([2dj\bar{\sigma} + h] \beta) \quad (8)$$

Now comes the second crucial step in the mean field approximation: we say the $\bar{\bar{\sigma}}$ must be equal to $\bar{\sigma}$. This of course would be true if we had a big enough sample and we were deep in the interior of it (and, as said, d was high). This *self-consistency* trick leads to

$$\bar{\sigma} = \tanh ([2dj\bar{\sigma} + h] \beta) \quad (9)$$

That is our equation which determines $\bar{\sigma}$. The factors 2, d , j and h are just numbers, so it is an equation in $\bar{\sigma}$.

Let's just remind ourselves what the solution looks like. We did another small manipulation. We defined the new variable

⁷Remember that the hyperbolic tangent function comes from working with the partition function Z of the spin, given by equation (3). Then we used the familiar fact that the average energy of a system – here just one spin – is given by $-\partial \log Z / \partial \beta$. In the course of the calculation appeared first the cosh function and then the tanh function.

$$y = 2dj\bar{\sigma}\beta \quad (10)$$

And we wound up with the equation

$$\frac{y}{2dj\beta} = \tanh(y + \beta h) \quad (11)$$

Now since $\beta = 1/T$, this is also

$$\frac{Ty}{2dj} = \tanh\left(y + \frac{h}{T}\right) \quad (12)$$

First exercise: what happens if $h = 0$? Then it is simply $Ty/2dj = \tanh y$. And now, just by inspection, we get a pretty good idea of what the solution looks like. On the right side, $\tanh y$ looks like this, figure 4.




Figure 4: Graph of $\tanh y$.

And on the left side we have a linear function of y which passes through the origin. It has the slope $T/2dj$, where T is the temperature. So, in particular, when the temperature is high, the slope is steep. And the only place where

these curves intersect is right at the origin, figure 5.




Figure 5: Intersection between $Ty/2dj$ and $\tanh y$ when the slope of the straight line is steep.

Therefore, for large temperature, there is no magnetization. Indeed the solution to equation (12), when $h = 0$, is $y = 0$. But y is proportional to $\bar{\sigma}$. Therefore $\bar{\sigma} = 0$ when the temperature is high, as expected.

Now we start lowering the temperature. The straight line rotates clockwise while still going through the origin. Its slope decreases. When T reaches $2dj$ it becomes tangent to the hyperbolic tangent curve at the origin.

We saw all that in chapter 9. We are only going through a brief review, before turning to liquid-gas transition, because sometimes a second presentation may clarify points which had remained unclear or which we had missed.

So that means, at temperature $T = 2dj$, when the slope of the straight line is 1, which is also the slope of \tanh at the origin, that it is the magic point where the two curves lie on top of each other.

Go past that. Decrease the temperature even more. That is, decrease the slope of the straight line even more. And now we have three solutions to equation (12). There is still mathematically the solution at the origin. But now there are also two more solutions, with $y > 0$ and $y < 0$, figure 6.




Figure 6: Intersections between $Ty/2dj$ and $\tanh y$ when the slope of the straight line is less than 1.

Those solutions where $y \neq 0$ are magnetized. Indeed, $\bar{\sigma}$ being equal to $y/2dj\beta$, see equation (10), it is non zero either.

That is a phase transition. Past that point, when now $T < 2dj$, the system doesn't quite know what to do. It doesn't stay at origin because it is unstable. But should it go to the solution point with $y > 0$, or should it go to solution point with $y < 0$?

In other words, there is an ambiguity. It seems like the spin at the center of figure 3 wants to have a non zero average value. But is that average value up, or is it down? On the right side of figure 6 the average value of the spin is up, $\bar{\sigma} > 0$. On the left side it is down.

And the problem is that there is a complete symmetry in the situation. If you take all the spins and, whatever each one is, turn it over, you end up with the same situation and therefore the same problem back again.

So there is no way to decide whether the spin will prefer to point up or to point down, that is whether it will choose the solution with $y > 0$ in figure 6, or the solution with $y < 0$.

To decide we add a tiny little bit of magnetic field h . The effect of that is to shift the tanh curve a tiny bit either to the left or to the right. But it is enough so that the point $y = 0$ is no longer a possible solution, even unstable. And only one of the two where $y \neq 0$ will remain, depending on the orientation of the tiny extra magnetic field, because one solution will have much less energy than the other.

That is what happens. That is magnetization. If the extra tiny external magnetization, however small, pushes the spins to be up on average, because that is where they have less energy, see equation (4), then up will win. If the extra magnetization pushes them to be down, then down will win.

We shall draw a graph of what the magnetization of the system looks like as a function of two variables: the temperature and the external magnetic field. Those are the two parameters that can vary in the problem. We can turn the temperature up or turn it down. We do that by heating up the heat bath or cooling it down, with a little fire under it or with cold air. And we can change the magnetic field by running a current through a coil, and make the magnetic field larger or smaller, in the direction we like.

What about j ? The parameter j is also something which we could change. In order to change it what we usually have to do is change the material. In fact we can also change j by changing the pressure or other things. But we are going to consider j to be a fixed number. It is fixed for the material has a fixed chemical composition, at a fixed pressure, fixed distance between the atoms, etc. The parameter j is a well-defined number which we can look up in a chemistry handbook. But β and h are the parameters that are easily under our control.

So we might want to ask: as an experimental fact, what does the magnetization look like as a function of temperature and magnetic field?

The temperature $T = 2dj$, where a transition appears, is called the *critical temperature*. It is not exact. It is a mean field approximation. It becomes exact when d becomes large. For $d = 2$ it is off by a significant amount. For $d = 3$ it is already very good. But in any case it captures qualitatively what really is happening⁸.

So let's denote the critical temperature T_c , read T *critical*. And we are in a lattice of dimension $d \geq 2$.

Below T_c , the system is spontaneously magnetized. Everybody wants to line up. The lining up tendency wins, and everybody lines up that way. Not absolutely. They are not a 100% lined up. It fluctuates but there is a net, in the sense of average, tendency for everybody to line up.

⁸Remember that for $d = 1$ we saw that there is no phase transition, no critical temperature.

Above T_c , *thermal fluctuations win*. And the system doesn't magnetize.

When we talk strictly about real magnets, T_c is also called the Curie temperature⁹. Here, however, we are talking about the general phenomenon of phase transition. So it is called the critical temperature. The Curie temperature is a special case specific to real magnets. The "lambda point", in liquid helium, is another special case.

They all fall into a class of transitions of things which are very similar in many respects. We will see why, or at least how the transition in magnets is related to another interesting transition.

Now let's draw a picture to study the magnetization as a function of T and h , figure 7.




Figure 7: Study of magnetization in the (T, h) plane.

⁹Named after Pierre Curie (1859 - 1906), French physicist, famous also for being the husband of Marie Curie (1867 - 1934).

We positioned T_c on the T axis, because it is the critical temperature, defined when there is no external magnetic field, i.e. $h = 0$, and at which things happen.

First of all, if we are at a higher temperature than T_c , and there is no magnetic field, $h = 0$, then the magnetization is just $\bar{\sigma} = 0$. There is no bias due to an external magnetic field. At high temperatures the statistics of ups and downs will equally win. And so, on the T -axis, on the right of T_c , $\bar{\sigma} = 0$.

Now, sliding along the T -axis starting from a point farther than T_c , we begin lowering the temperature, that is we move leftward. When we get to the critical point, the system just doesn't know what to do. Should it go up or should it go down?




Figure 8: Magnetization on or close to the T -axis.

The tiniest little bit of stray magnetic field h will tell us which way to go. If h is positive the magnetization will be positive. Right above the T -axis it won't be zero, and right

below it won't be zero either. Spontaneous magnetization has happened, figure 8.

It corresponds to the two possible roots, with the straight line in figure 6 either intersected on the right side or intersected on the left side. Physically, it is just that the extra little bit of magnetic field has broken the symmetry between ups and downs and has biased the sample. That tiny little bit of biasing h is enough to make $\bar{\sigma} > 0$ if h points up, and $\bar{\sigma} < 0$ if h points down. See equation (2) for the effect of the sign of h .

Now we can ask : what is the jump in the magnetization going across the T -axis? Remember that $\bar{\sigma}$ is not necessarily $+1$ or -1 . $\bar{\sigma}$ is the *average* magnetization. It can be 0 or it can be $2/3$. It can be $-1/4$ or any number between -1 and $+1$.

At zero temperature, i.e. at the origin, the absolute value of $\bar{\sigma}$ doesn't depend on the magnetic field.




Figure 9: Magnetization along the h -axis.

When $T = 0$, all the spins are lined up. $\bar{\sigma}$ is either $+1$ or -1 . Where $h < 0$, all the spins point downward, therefore $\bar{\sigma} = -1$. And where $h > 0$, they all point upward, and $\bar{\sigma} = +1$, see figure 9. That comes from the fact that at zero temperature the spins all choose the lowest energy, and there is no fluctuation.

What about life at the critical point? Imagine going up and down from $h > 0$ to $h < 0$ at $T = T_c$. At exactly the critical point, the straight line in figure 6 has exactly the slope 45° . So at the critical point, the magnetization is clearly 0.

Now, from T_c , we decrease the temperature a little bit, while staying on $h = 0$. That is, we don't turn the external field. We lower the slope of the straight line in figure 6. The two branches of $\tanh y$, the one on the right and the one on the left, now play a role. As we saw, the magnetization $\bar{\sigma}$ doesn't stay 0, because it is unstable. The reason is because *there is always a stray magnetic field h , however tiny, the consequence of which is that $\bar{\sigma}$ becomes either positive or negative.*

Of course, as we move vertically upward, staying just to the left near the critical temperature, the magnetization increases. But it does so continuously in the sense that, in the vicinity of the point $(h = 0, T = T_c)$, to the left of it, when $h > 0$, $\bar{\sigma}$ is positive but very small, and, below the T -axis, $\bar{\sigma}$ is negative but very small.

In summary, as we come in to the critical point $(h = 0, T = T_c)$, the magnetization tends to 0 from every direction, including arriving from the left.

However, once we are away to the left of the critical point, there is a jump in the magnetization. There is a shear in the magnetization surface $\bar{\sigma}(h, T)$ along the half line $h = 0, T < T_c$. But, as just explained, even that shear goes to zero as we approach T_c .

Let's pursue our investigation of the surface $\bar{\sigma}(h, T)$ in some detail at different points. For example, let's look first at a temperature above the critical one. If we had our external magnet turned on, and we slowly varied the magnetic field h from a positive value to zero, while staying at a fixed temperature T above the critical temperature T_c , then the magnetization $\bar{\sigma}$ would get smaller and smaller, going to 0 continuously, figure 10.




Figure 10: Lowering h while staying at $T > T_c$.

And if we passed into $h < 0$, $\bar{\sigma}$ would increase again in the negative values (that is $|\bar{\sigma}|$ would increase, while $\bar{\sigma} < 0$). And all this would happen completely smoothly, just as we

might expect.

Now let's look at a temperature below the critical one, and for the sake of variety let's position ourselves in the $h < 0$ half plane¹⁰. As the experimental physicists manipulate the magnetic field at fixed temperature, the magnetic field varies. At first the magnetization is nice and smooth until it reaches the T -axis, figure 11.

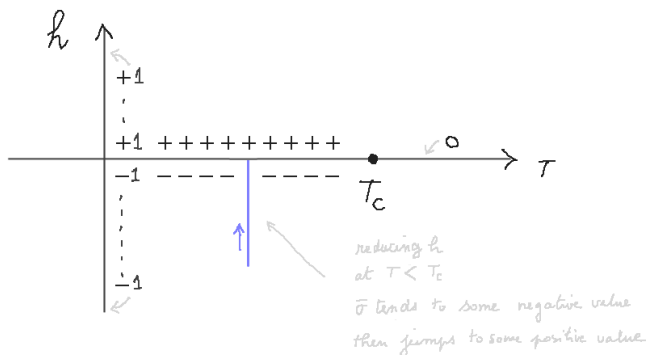


Figure 11: Varying h while staying at $T < T_c$.

The limit at $h = 0$ coming from negative values, is some magnetization $\bar{\sigma}$ negative. If we are not afraid of heavy notation, we can denote it $\bar{\sigma}_-(T)$. Then change a tiny bit the magnetic field h making it positive: the magnetization changes discontinuously to some positive value, $\bar{\sigma}_+(T)$. The jump is biggest at the origin, and goes to zero as the temperature approaches T_c .

¹⁰Actually it is a quadrant, because T is always greater than or equal to zero.

Phase transitions are classified by their order, which depends on how smoothly the free energy varies with respect to the parameter. Technically what happens in figure 10, when we cross the T -axis somewhere between 0 and T_c , is a second-order phase transition. But we don't need to distinguish them because we are not going to get a chance to talk about first-order phase transitions.

In summary that is the jump which takes place on the segment where $h = 0$ and $0 \leq T \leq T_c$, i.e. where there is a shear in the surface $\bar{\sigma}(h, T)$.

This jump in magnetization from $\bar{\sigma}_-(T)$ to $\bar{\sigma}_+(T)$, when we cross the T -axis at some temperature T below the critical one, is a symptom that in a sense there are two phases.

Suppose we didn't know anything about what happens on the right part of figure 11. Suppose the critical temperature was too high for our experimental physicist to ever study and become aware that things behaved differently beyond it. Then he would say that there are two phases. There is a phase of negative magnetization; there is a phase of positive magnetization; and, as we vary the parameter, it jumps.

It is not that all the spins flip from down to up. It is the average that goes from some value down to some value up. And, as we stressed, the jump is biggest near $T = 0$ and goes to zero as T approaches T_c .

Now there is another way we can move around the space (h, T) . We might wonder: is there a way to move from the positive zone to the negative zone without a sharp change?

And of course there is. We can go *around the critical point*. Then we experience no jump. But we do get from positive magnetization to negative magnetization.

So, on the one hand, experimentalists starting from a point (h, T) in the lower part, and left of T_c , and who have never heard about the critical point, will say: "there are two phases. They are distinct. And not only are they distinct, but there is a sudden jump as we vary the magnetic field h ."

On the other hand, someone with more experience with the system would say: "oh, no, no, no. There is really not a sharp difference between the two zones. Let me show you how to go from one to the other without jumping."

And both would be right. Both statements are correct statements. There is a jump if we stay at fixed temperature below T_c . But if we allow ourself to explore the whole region of (h, T) , there are ways to go from the upper zone to the lower zone which don't involve the jump, which are smooth.

Now we have everything we need to know about magnets. But we are interested in the liquid-gas phase transition.

Liquid-gas phase transition

What does the above study of the surface $\bar{\sigma}(h, T)$ have to do with liquid-gas transition? There is a way to think about exactly the same system as being the transition between liquid and gas. It is not good for liquid to solid. This latter

case is very different. So let's talk about the former and what it possibly can have to do with a bunch of particles.

Consider a box of particles. It can be gas or liquid, so let's say a box of fluid, And let's plot its potential energy, figure 12.

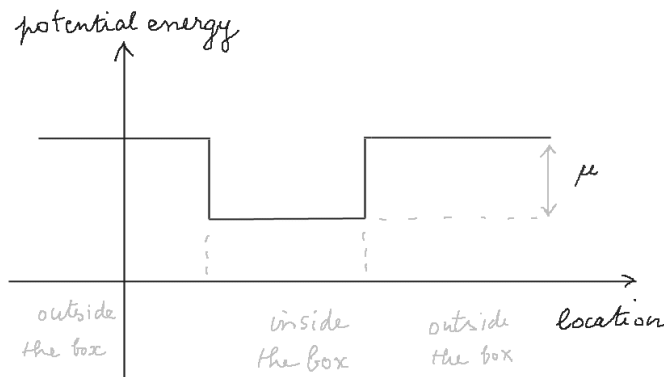


Figure 12: Potential energy of a box of fluid.

The vertical edges correspond to the walls of the box¹¹. In between, we are inside the box. Elsewhere, we are outside.

If it turns out the box really corresponds to some macroscopic material, or whatever it happens to correspond to, then the energy of a particle is lower when it is inside the box than it is when it is outside the box.

Now we could manufacture this system by having the inside of the box and the outside of the box being of slightly

¹¹They *don't* represent the walls though. This is only a graph of the potential energy.

different elevation. We could have a flat box with a lower bottom over some area. But for whatever reason we are going to imagine that a particle at rest inside the box has less energy than relative to what it has outside the box.

That difference in energy has a name. It is called *chemical potential*. And it is usually written μ . It is actually the amount of energy that it takes to remove the particle from the box.

In this particular case it is not important to know where the difference comes from, why are there different energies. We just take it that the energy is lower inside the box.

Furthermore, there are particles inside the box and outside the box. They are all over the place, figure 13.

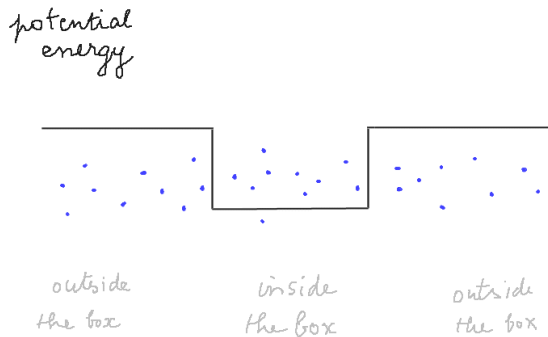


Figure 13: Particles inside and outside the box.

One more thing: the walls of the box are permeable. Particles can come in and go out. They can cross the boundaries

of the box.

This could be, for instance, a bunch of molecules in the following system. We dig a hole in the ground. And we deposit some particles in as well as out. If the particles are in the hole, they have less potential energy than if they are out of the hole.

We can ask what is the density inside the box and outside the box. We are thinking of thermal equilibrium. Where is the density of particles likely to be larger, in the hole or out of the hole? In the hole.

Indeed, the energy is less in the hole. And so it is natural that the density there be higher, because the Boltzmann distribution always favours lower energy.

Remember, the Boltzmann distribution is $P_i = e^{-\beta E_i} / Z$, where P_i is the probability of the state i , which has energy E_i ¹². The function – that is the probability – decreases as the energy increases. So it always prefers lower energy.

So the answer will be that the system will adjust itself so that the density in the box is in a certain kind of equilibrium. And the equilibrium will depend on the energy cost

¹²That is the Boltzmann distribution – at thermal equilibrium by definition – in the discrete case, which is often easier to conceive. Remember, there is a big space of states Ω . The individual possible states, which are the elements of Ω , are labelled ω_i . The i -th state has energy E_i . And the probability that the system be in the i -th state is P_i . We constructed the whole model in chapters 1 to 4. The reader is invited to refresh his or her memory going back over them. In view of all we have done since, they may make more sense, and be easier to follow.

of taking a particle outside the box.

Now we can set this up in a very definite way and do ordinary statistical mechanics. There is a potential energy which depends on position. But it depends on position in a very special way. It is flat inside of the box of interest, that is between the walls in figure 13. And it jumps when we go outside the box.

We focus on inside the box and ask the question: what is the density of the fluid as a function of the depth of the chemical potential μ , see figure 13?

Note that μ is called the chemical potential because in general it is different for different kinds of molecules. So we could imagine a situation where one kind of molecule prefers to be in the box, and the other kind of molecule doesn't prefer to be in the box. Turning on the chemical potential then can change chemical composition in various ways.

However we are not interested in the chemistry here. We are interested just in this idea that we can tune the density by changing the chemical potential. All that μ is here is a term in the energy which depends on the number of particles in the box. There is a factor μ for each particle. If there are N particles in the box, then $N\mu$.

Now in this problem the number of particles is a variable, because particles can come in the box and go out of the box.

That means a configuration depends not only on the position and momentum of every particle inside the box, but also on the number of particles. That number is free to

change because particles move in and out of the system. And part of determining thermal equilibrium is determining how many particles are on average inside the box.

We are not going to go through this in detail, but we will outline the calculations. The usual formula for the partition function, e to the minus β times the energy of each state, properly summed, will in this case look like this

$$Z = \sum_{PXN} e^{-\beta E + N\mu\beta} \quad (13)$$

The term E in the exponent represents the energy of one state. As usual it is multiplied by $-\beta$. But now to it is added the term $N\mu\beta$, which is the potential energy of that state. There is a plus sign because the potential energy of one particle inside the box is $-\mu$ compared to one outside. And there are N of them. Finally the sum is over all configurations (X, P) and all N .

We are not going to work it out, it would make the chapter too lengthy, but we shall see what the meaning of the chemical potential is. The factor μ is the thing that we adjust if we want to adjust the density of the fluid.

Actually we could do a lot of things to adjust the density. We could seal the box and not let the particles go in and out. And we would put in just the number of particles we wanted.

But we may be interested in the problem where particles can go freely in and out of the box. In that case, we change the density by changing the chemical potential. By varying

it, we can vary the average number of particles in the box.

In this language what is a liquid-gas transition?

We take our box, we keep it at a fixed temperature, and we start varying the chemical potential. The more we increase the chemical potential, the more it wants to pull particles back into the box. Of course it depends on the size of the box, but there is a density that is a natural function of the chemical potential.

And how does the transition materialize? As we vary the chemical potential, at some point the density of the fluid suddenly changes. The fluid changes from being a gas to being a liquid.

The box, suggested in figure 13, of course might be a region of three-dimensional space. We can also imagine a real three-dimensional box. When a particle goes in, we get actually energy out of it, because of the chemical potential difference. And when a particle goes out, it costs energy. The extra energy released, which is the same for every particle going into the box, will tend to increase the density inside the box.

Even though the atmosphere around the box might be at a temperature and density corresponding to vapour, if we keep lowering the energy inside the box, sucking particles in, then all of a sudden we hit a transition point where everything becomes liquid. The difference between liquid and vapour is some discontinuity in the density, liquid being denser than vapor.

So if we could do this experiment with the box, vary the potential energy of particles inside the box, while outside the conditions were such that we had steam vapor, particles would congregate inside the box as we lower their energy. And all of a sudden at some point we would jump to the liquid phase.

That liquid-gas transition is very much like the earlier transition studied in the previous section, see for instance figure 11. Let's spell out exactly what the relationship is.

First of all, what are the conditions required to have a liquid-gas phase transition? Two conditions need to be satisfied:

1. There must be a hard core repulsion between molecules. In fact, molecules always have hard core repulsion. You cannot stick two molecules into the same site. For a number of reasons, they will repel. So it is natural to view the molecules of our fluid like little billiard balls. They don't want to interpenetrate. Roughly speaking you cannot put two in the same site. And that is a form of really hard core repulsion.
2. The next thing important is that there be a little bit of attraction between molecules when there are not quite touching.

So we need a potential energy that is big, repulsive, positive when the molecules try to get into each other. And then, when we separate them a little bit, all of a sudden they want to attract a little bit. That is an extremely common

feature of molecular interaction: hard core repulsion and very short range attraction. That is what is needed to have the standard liquid-gas transition.

We shall see how the magnet produces exactly that situation. Let's start with the magnet with all of its little elementary magnets pointing down. When they are all down, that is the ground state¹³. And let's call that empty space.

In fact what we are going to say is something like this: on each site of the lattice we can have a particle or not. This is our aim. Out of our liquid gas problem we are going to make a lattice version of it.

The particles, or molecules, live on a lattice. At a site, as said, there can be a particle or no particle. When a site has no particle it is assimilated to *down*. And when all the sites are empty, that is the *ground state*, figure 14.




Figure 14: Fluid modeled as a lattice. Represented is the ground state with all the sites empty.

¹³This is a change from equation (2), where the spin lower energy was up. It is entirely conventional. We just have to think now of h being a negative number, or in the equation change the sign in front of it.

So in terms of σ at a site we can have

- $\sigma = -1$: no particle at the site
- $\sigma = +1$: one particle at the site

What about two particles at a site? Well, the magnet doesn't allow that. The value of σ can't be 2. It can be only -1 or 1 . It cannot be 4. It cannot be anything else.

Thus, immediately the model forbids the possibility of two particles on a site. And that is exactly what we want. We are trying to model a gas where we cannot put two particles on a site because molecules have a hard core. Automatically built into the system of particles, if they are described by the same kind of variables $\sigma = \pm 1$ as the magnet, they have an infinite hard core repulsive barrier when they try to get onto the same site. So we satisfy above condition number one.

What about condition number two? Do they attract? Here we are going to use for the energy exactly the Ising model of energy plus a little magnetic field. And we shall see how this translates in the language of our particles of fluid which are hard core *and* have short range attraction.

In fact we can keep on using the magnets simile or we can use the particles simile. The point is it is the same mathematical system. So let's not worry about it, and start with the term already given by equation (1)

$$\sum_{links} -j\sigma_n\sigma_m \quad (14)$$

It is the energy contained in each pair of neighboring sites, summed over all the links.

How much energy is there if there are no particles whatsoever? That is the situation where all the σ 's are down. Then there is an energy $-j$ for each link.

But this total ground state's energy, which is $-j$ times the number of links, what do we do with it? We don't do anything with it. That is just the energy of the system when there is no particles at all. Does it matter to anything? No, it doesn't. It is always there. It is the ground-state energy. And we can throw it away.

We could get rid of it by subtracting $-2j$ times the number of links. We would get

$$\sum_{links} -j\sigma_n\sigma_m + 2jN_{links} \quad (15)$$

But we might as well just ignore it. It is just a number. And the energy differences are the important thing. So we will continue to use expression (14) as our energy function when there are no particles. That gives us a certain energy. To nail the point – and use the prerogatives of physicists:-) – let's just call it zero. In other words, we set the zero of energy when there are no particles.

Now, what happens if we put in one particle? That means flip one of the spins over, figure 15.




Figure 15: One site occupied by a particle. It is equivalent to one spin flipped to positive.

All the other spins are still down. But one is up.

Again, to flip one spin up or to introduce one particle are mathematically equivalent operations. So the reader should not be disturbed if we go back and forth between the two ways to describe what is going on in the model.

Let's calculate the energy that we get. How many bonds have we broken? It depends on the dimensionality. Let's stick to the model in 2D because it is easy to visualize. If we flip one site over, then we break four bonds, figure 16.

The amount of energy necessary to break one bond is $2j$. That yields $8j$ for all four bonds in the figure. Later on when we add a little magnetic field, we will add $2h$. But for the time being, $8j$ is the energy it costs to create one particle at the point shown in figure 16.




Figure 16: Bonds broken from flipping one site over.

What about two particles? Suppose they are not adjacent. We don't need to draw another figure. It is easy to see that, in that case, it costs $16j$ to put in the two particles – even if they are separated by only one diagonal interval.

And if they are neighbors, if they are adjacent? When we introduce 2 adjacent particles, we break only 6 bonds. The one between them is no longer broken, because both sites have flipped. So in terms of energy we pay only $12j$.

If we plotted the energy as a function of the distance between particles, we would find that just the mere act of putting two particles in the lattice, farther apart than one link, costs us $16j$. But that is alright; that energy is always there, no matter where we move the particles around.

However, if we move the two particles within a bond length, the energy decreases. That is like a potential energy which decreases when they get very close together.

So there is a potential energy, which depends on the posi-

tions of the particles. And when the particles get within one bond length of each other, the energy decreases by $4j$.

That is like having a short-range attractive energy. When the two particles are close to each other, the energy is negative relative to what it would be if they were far apart. So particles attract. They like to be close together.

How close together? One bond length away. But that is not an overwhelming attraction. It doesn't say that the energy is infinitely negative if you put them next to each other¹⁴. There is just a modest amount of attraction, a modest saving of energy of putting them close together. And it is very much like the molecular attraction of a pair of molecules when they are close together.

So, we have a system of elements now (σ 's, spins, small magnets, particles, whatever) which is mathematically isomorphic to a system of hard core particles on a lattice, which have a short-range attractive force between them, and where we could have any number of particles.

The particles number is something that can change. So it is like a system of molecules where they can come in and go out of the system. It doesn't have a definite number of molecules. Their number itself is a variable.

There is short-range growing negative potential, that is short-range attraction. Yet, there is very short-range infinite repulsion. We cannot put two elements on the same

¹⁴That is the case of attraction laws in $1/r^2$. Then the delta potential energy is proportional to $-1/r$. So it is infinitely negative for instance when point masses are at distance zero of each other.

site. But when they get a little bit apart, they are slightly attractive.

That is exactly what we need for liquid-gas phase transition. Remember the two conditions spelled out above.

Now let's add something else. By the way, does this system have a chemical potential? In a chemical potential the energy is stored in just having one particle. That is what it is. The chemical potential is the energy in having a particle present. Just by virtue of having the particle, if there is an energy that wouldn't be there if the particle weren't there, that is called the chemical potential.

So the answer is yes. Just having a particle, with no other particle around, gives us an energy $8j$. Therefore there is a chemical potential.

But we want to be able to vary that chemical potential, without varying j incidentally. We would like to hold j fixed, and vary the chemical potential separately from it.

That is easy to do. Put a magnetic field in. From equation (14), this takes us back to equation (2). But, since we want the energy to be lower when σ_n is negative now, i.e. when there is no particle, we modify it a little bit. The sign in front of the second sum becomes plus. The energy of the whole system is

$$E = -j \sum_{links} \sigma_n \sigma_m + \sum_{sites} h \sigma_n \quad (16)$$

How much does this give us when we add a particle to the brew? At the n -th site, σ_n starts down, that is -1 ,

when there is no particle. When we introduce a particle, σ_n switches to $+1$. So it jumps two units. It gives us two units worth of energy for every spin which is flipped up, in other words for every particle that we put in. That is $2h$.

So for each isolated particle, we can add $+2h$. The total energy needed to introduce an isolated particle is

$$8j + 2h \tag{17}$$

Now we have everything we need to study the liquid-gas transition in a fluid of molecules. From our Ising magnet model we have built a system that is equivalent to a collection of particles on a lattice. It has the following three specifications:

1. It has a variable chemical potential. Notice how we vary it. We do it by varying the magnetic field. The magnetic field gave us an energy per particle.
2. It has the hard core billiard ball potential which is repulsive.
3. And it has a short-range attraction, the coefficient of which is $4j$ (going from $16j$ when the two particles are far apart, to $12j$ when they are separated by only one link).

We want to keep the molecular properties fixed, therefore we are not going to play with j . That has to do with the molecular properties. By design, it tells us what the potential energy between them is.

All we can vary is the chemical potential. And that is the same thing in our model as varying h .

So the problem of the behavior of the molecules of a fluid, with the three specs listed above, has become exactly the same as the previous problem of a magnet plunged into an external field, which was summarized in figure 11. They obey precisely the same model.

There is a temperature that we can vary. Particles are at some temperature. And there is a chemical potential that we can vary. Now h is not exactly the chemical potential. There is an offset by an amount $8j$.

What about the particle density? We haven't discussed it yet. So let's do it. How many particles are there, on the average, at a particular point? The answer is somewhere, first of all, between 0 and 1. Of course the number of particle is either 0 or 1, but we are talking about the average.

The average number of particles at a site is

$$\frac{1 + \bar{\sigma}}{2} \tag{18}$$

Let's see why that is true. The number of particles itself is the random variable $(1 + \sigma)/2$. Indeed, if there is a particle there, that means $\sigma = +1$. One plus one is two, divided by two is one. So the r.v. has the right value. And if there is no particle? $\sigma = -1$. In that case $(1 + \sigma)/2$ is 0. It is ok too. And the average of $(1 + \sigma)/2$ is $(1 + \bar{\sigma})/2$.

Since the density ρ of particles¹⁵ is proportional to the av-

¹⁵Here we talk about the average density, not about an r.v.

erage number of particles, with units properly chosen, we can write

$$\rho = \frac{1 + \bar{\sigma}}{2} \quad (19)$$

That is the density. Notice that it is not $\bar{\sigma}$ itself, but is given by formula (19). There is no subtle fact here. It is just that we have naturally chosen our spins to take their values in $\{-1, +1\}$, rather than in $\{0, 1\}$. So, to produce 0 or 1 from the two possible values of σ , we have to use the function $(1 + \sigma)/2$.

So the density of the fluid is just $1/2$ plus half the average magnetization.

Now we are all set to see what would happen if we were to vary the chemical potential, or equivalently h , while keeping the temperature fixed. Instead of following the evolution of the magnetization $\bar{\sigma}$, we will follow the evolution of the fluid density $\rho = (1 + \bar{\sigma})/2$, figure 17.

If we are above the critical temperature, the density ρ varies continuously, with no jumps. In a box of gas, at fixed temperature above T_c , as we decrease the chemical potential, i.e. as we lower the energy, as we pull it down, nothing very exciting happens when we cross the horizontal axis in figure 17. The horizontal axis corresponds to $h = 0$, or to a chemical potential of $8j$, because it has the general form $8j + 2h$.




Figure 17: Space of variation of the fluid density $\rho(h, T)$. This is exactly the same picture as in figure 7, except that the quantity $\bar{\sigma}$ we were looking at is now replaced by $\rho = (1 + \bar{\sigma})/2$.

Below the critical point, by contrast, as we vary the chemical potential h , all of a sudden, when we cross $h = 0$, the density jumps. Reducing h reduces the density, see equation (16). And crossing the horizontal axis is the phenomenon of boiling.

For a long time, experimentalists had difficulties experimenting with the chemical potential at temperatures above T_c . So they only knew of phase transition. What is the critical temperature for water? It is not 100°C . It is about 374°C . But the chemical potential is linked to the density, therefore it is pressure dependent. And it is hard to experiment above T_c .

Below the critical temperature, when we cross $h = 0$ from positive to negative is when boiling happens. And going the other way, it is liquefaction.

Above the critical temperature, when we cross $h = 0$ no boiling nor liquefaction happen. There is a smooth continuous transition of the density.

Questions / answers session

Q.: It is easy to understand how we change the magnetic potential. But how do we change the *chemical potential*?

A.: Really by varying the *density*.

We do whatever it takes to vary the density of the fluid. For example we could create a difference of potential in energy between outside the box and inside the box, and we would literally suck particles in. We would have to devise something in practice. And it would not be very accurate.

Actually that is not the way we do it in practice. What we do instead is vary the *pressure*, figure 18.

Varying the pressure is the simplest solution to vary the chemical potential. When we do it while the temperature remains below the critical one, we witness a sudden jump in the density of the fluid. On the left of T_c , there is a frontier between liquid and gas. Below it is gas, above it is liquid.




Figure 18: *Simplified* temperature/pressure diagram. The solid phase and the "triple point" are not shown. But we see the frontier between the liquid and gaseous phases. The dotted path shows how we can go from one to the other continuously, with no jump, by going around the "critical point".

But notice, as before, that we can also go from gas to liquid, and back, without any jump, by going around the critical point.

Q.: When you say you are changing the chemical potential by changing the density, are you basically changing the effective distance for Van der Waals potential?

A.: Well, we don't change the properties of the molecules. Those stay the same.

But when we increase the density, we do shorten the average distance between molecules. At some point the system realizes, so to speak, that a negative potential energy is there.

Then it tries to pull all the molecules together. That is when it suddenly makes the transition.

But this description is valid only if we stay below the critical temperature. If we go around it on its right we can go from the liquid phase to the gaseous phase without jump.

Q.: In a liquid-gas transition, what is the relationship between the critical temperature and the boiling point?

A.: The boiling point depends on the pressure. This in turn means it depends on the density.

In figure 17, we cannot see it, because the coordinates are T and h , and the frontier is flat and horizontal.

But the pressure is a function of the temperature and the chemical potential. When we represent the average "magnetization" $\bar{\sigma}$ in the T and P coordinates, the frontier is no longer a flat interval along the horizontal axis. It is the curve shown in figure 18.

Then, for each pressure below the ordinate of the critical point, we can read off at what temperature the transition from liquid to gas happens. The lower the pressure, the lower the boiling point. On the top of Mount Everest water boils at 68°C .

One of the very fascinating things is the properties of the system near the critical point. There is a whole theory on the behavior of the system and there are lots of experiments

to study what happens as we approach the critical point.

The phenomena happening there are all characterized by what is called *critical exponents*. Every quantity that you can think of, that is interesting, in its dependence on temperature will generally go as

$$(T - T_c)^\gamma \tag{20}$$

where γ is some power. And those powers are not 1's and 2's. They are various odd numbers, rational, irrational or only experimental.

Of course they vary from one kind of phase transition to another, but they are exactly the same for the magnetic transitions and the liquid-gas transitions. They fall into exactly the same class.

The behaviors at the critical points are rather insensitive to the details. They depend on features that don't care whether we put the particles on a lattice. They don't care whether it is nearest neighbor or second nearest neighbor. They all behave the same way. And the magnet and the fluid transitions are in the same class.

Q.: Do the critical points form a continuum, or is it a discrete finite or countable collection of points?

A.: The set of critical points is discrete, but it is not finite.

Q.: How do you define gas versus liquid? If when we go around the critical point, there is no transition. So it seems that they should be the same.

A.: Indeed there is no sharply defined difference.

The density is obviously different in the lower part of figure 18 than in the upper part. If we try to go across the transition frontier left of T_c , the density jumps. If we go around the critical point, the density varies continuously.

The density of the fluid in figures 17 or 18 behaves like the magnetization in figure 11. The precise connection is $\rho = (1 + \bar{\sigma})/2$.

On the left of T_c there is a frontier where it jumps. But we can go from a point below the frontier to a point above it by going around the critical point, i.e. passing to the right of T_c , then there is no jump.

Q.: So liquid and gas coexist if we go around?

A.: Well, they are not distinguishable. There is a continuous transition from one to the other.

Q.: What do we need to add into the theory to get liquid and gas together in a phase?

A.: We need for instance gravitational potential energy to

separate them.

Q.: Could it also be surface tension?

A.: No. In most of these kinds of analyses as above we think of an infinite volume of fluid. So surface tension is not relevant here.

In an infinite volume of fluid, however, a gravitational field would create a liquid-gas physical separation. They would be separated by a plane, as we are familiar with.

There would be an altitude-related chemical potential. So varying the altitude varies the chemical potential. And there is a transition at some height.

Q.: In this model, do we get a heat of vaporization as we go across that transition?

A.: In general yes. No doubt.

The problem is that, in the magnetization model, the energy near the transition frontier is exactly the same just above and just below it.

So it is a good question. In the fluid case some refinements, at least in the analysis, and perhaps in the model, would have to be introduced.

Q.: Does this have implications for superconductors and magnets, in particular for highly sensitive detectors, for example detectors capable of revealing the presence of dark matter?

A.: A material going from ordinary conductor to superconductor does undergo a phase transition. It has similarities with what we studied.

Now concerning detectors, if you can adjust your superconductor to be very near the phase transition, then by a tiny change in some external parameter you can have it jump from superconductor to non-superconductor. And that gives you the possibility of constructing highly sensitive detectors.

For instance some energy is deposited in the detector. And even a small amount of energy can create a little local bubble of changed phase. So it can be useful as a very sensitive kind of detector.

That is true in fact for any phase transition. Think of cloud chambers for instance.

We are finished with the sixth course in the collection *The Theoretical Minimum* on statistical mechanics. And that completes the collection. These notes come from the site of the notetaker at https://lapasserelle.com/statistical_mechanics