

Lesson 7: Entropy versus reversibility

Notes from Prof. Susskind video lectures publicly available on YouTube

Introduction

In this lesson, we will eventually get to the second law of thermodynamics, what the second law means, how it is consistent. The subject is complex. To complete the main text, there will be several questions / answers sessions.

Before reaching the second law, however, we shall pursue the study of some physical examples which we have begun to examine with the tools of statistical mechanics.

First, we shall illustrate the physics of gases with some numerical data. We shall study the speed of sound in a gas, and in particular in a fairly dilute gas.

As a second preliminary example, we shall study the behavior of a harmonic oscillator in a heat bath.

Speed of sound in a gas

What is the speed of sound? There are two formulas for the speed of sound. But if you think about it for a moment, and you start with a very dilute gas, what would you expect the speed of sound to be?

Let's say we make a little bit of overdense region. How fast does that overdensity spread? That is the speed of sound. What would you guess? Well, it sure isn't going to be faster than the time that it takes for a molecule to travel out to a certain distance.

In fact, for a very dilute gas, we might expect that that little overdense region just spreads out with a velocity that is not too different than the average velocity of the molecules.

How does it spread out? It spreads out by the molecules moving out. How fast do the molecules move out? They move out with their velocity. What is their velocity? Now, we have to work a little bit to find out what the velocity is.

So let's write down a formula. In thermal equilibrium at a temperature T , in a dilute gas, every molecule has an energy

$$E_{part} = \frac{3}{2}k_B T \quad (1)$$

We are going to put the k_B in now, because we want to do some calculations in laboratory units. So the temperature is measured in Kelvin. Kelvin means relative to absolute zero, and 1 unit corresponds to 1 degree centigrade. That is 9/5 of a Fahrenheit degree.

From formula (1), we can see that the temperature goes to zero when the energy E_{part} of one molecule goes to zero.

That energy must be the kinetic energy of the molecule. It is the only energy that a molecule has for an ideal gas in a big box. So we have got to have

$$\frac{3}{2}k_B T = \frac{1}{2}mv^2 \quad (2)$$

where v is the average velocity of one molecule. So equivalently

$$v^2 = 3 \frac{k_B T}{m} \quad (3)$$

You might expect that this is more or less the speed of sound in a gas. And actually it is not too far off.

There is a more exact formula. It is established by actually studying the mechanics of lumps of gas and the forces on them. We are not going to do it in this lesson. Newton did it. It is not too hard.

You take a box of gas. You take a cubic region in the gas. And you ask: What forces are around it? The forces have to do with the pressure from the outside of the region¹. There is a gradient of pressure. There is some force due to the gradient of pressure. What is the response to the variation in pressure? It is $F = ma$, where a is the acceleration. So the box starts to move. The acceleration depends on the mass in the box. So the mass density comes into play. But in the end you find a formula.

Newton's formula is

$$C^2 = \frac{1}{m} \frac{\partial P}{\partial \rho} \quad (4)$$

We write C instead of v because this is a fairly general formula. It is more general than just for the ideal gas. C^2 is the partial derivative of the pressure with respect to the mass density. And since we used ρ to denote the *number* of particles per unit volume, the mass density is $m\rho$ where

¹Now we do continuum mechanics. The gas is taken to be a continuum.

m is the mass of one particle. Hence formula (4).

Let's see what formula (4) yields for the ideal gas. We know that for the ideal gas we have

$$PV = Nk_B T \quad (5)$$

or

$$P = \rho k_B T \quad (6)$$

The derivative of P with respect to ρ is

$$\frac{\partial P}{\partial \rho} = k_B T \quad (7)$$

And so, dividing by the mass of one particle, we get

$$C^2 = \frac{k_B T}{m} \quad (8)$$

It is not quite the same formula as formula (3).

The simplistic formula (3), when we take the square root to compute v , is too high by a factor of $\sqrt{3}$. That is approximately 1.7. So it is not so bad. In order of magnitude, it is correct.

Let's work it out for the air. We shall stick some numbers into formula (3). The Boltzmann constant is

$$k_B = 1.3 \times 10^{-23} \quad (9)$$

Air is mainly made of nitrogen. So we shall take for m the mass of a nitrogen molecule. Gas nitrogen has a diatomic

molecule. The atom of nitrogen has atomic number 7 and atomic weight 14. (There are 7 protons and 7 neutrons in the nucleus of the common isotope.) We shall take

$$\begin{aligned} m &= 28 \times \text{mass of a proton} \\ &= 28 \times 1.7 \times 10^{-27} \text{ kg} \\ &= 5 \times 10^{-26} \text{ kg} \end{aligned} \tag{10}$$

And we need to know room temperature.

$$T = 300 \text{ K} \tag{11}$$

in the same units where k_B as the value given above.

Using formula(3), we get

$$\begin{aligned} v^2 &= 3 \times 1.3 \times 10^{-23} \times 300 \times 0.2 \times 10^{26} \\ &= 23,4 \times 10^4 \end{aligned} \tag{12}$$

Therefore, in round numbers,

$$v = 500 \text{ m s}^{-1} \tag{13}$$

This is the correct value for the speed of the nitrogen molecules in the air at room temperature.

As an estimate of the value for the speed of sound, as expected, it is a bit on the high side. The right answer, in the air², is about 340 meters per second.

²Air is a mix of 78% dinitrogen, 21% dioxygen, and traces of other gases.

Questions / answers session

Q.: In assimilating formula (3) to the square of the speed of sound, we also assume that when a sound wave moves through the gas, the temperature is constant, don't we?

A.: Yes, we did assume that the temperature was constant. I think it is a fairly good assumption. For small amplitude sound waves, in a dilute gas, temperature doesn't vary very much as it vibrates.

I am not absolutely sure about it. But I would guess that the temperature has very little variation, in a near ideal gas, for very low amplitude sound waves. Then taking the temperature to be constant is probably ok.

Q.: According to this reasoning, would not the variation of pressure also be small?

A.: Yes the variation of pressure would also be small. But the variation of ρ too. And the rate of change in formula (4) – if that is what you are driving at – does not have necessarily to be small.

Q.: If we increase the pressure, won't we expect the speed of sound to be faster? Like in water or in a steel bar, it is much faster than in the air.

A.: Well the $\partial P / \partial \rho$, in formula (4), is calculated at the

pressure you are working at.

$\partial P/\partial\rho$ is not necessarily a constant. The pressure depends on the density in a complicated way.

In particular, as the molecules get really squeezed together, the derivative of pressure with respect to ρ gets very large.

So, you are right. But it is because $\partial P/\partial\rho$ changes.

Q.: Since the speed of sound is related to pressure, why doesn't pressure appear in formula (3)?

A.: In general, the speed of sound does depend on the pressure.

If you use formula (4), $\partial P/\partial\rho$ depends on P .

But if we are in the case of an ideal gas, the two following formulas hold

$$v^2 = 3 \frac{k_B T}{m} \quad (14)$$

where v is the speed of a molecule, and

$$P = \rho T \quad (15)$$

The calculation yielding $v = 500 \text{ m/s}$, was done for an ideal gas at $T = 300 \text{ K}$. Being in the case of an ideal gas, ρ and P are small. Moreover, at fixed temperature they are proportional to each other. As a consequence it eliminates

them from the formula for the speed of sound, either in its simplistic formula (14) or in the more accurate formula (8).

So we can state:

In the ideal gas range, the speed of sound depends only on the temperature.

In other words, in the ideal gas range, at fixed temperature, the speed of sound doesn't depend on the pressure or the density.

But, as we increase the density and get away from the ideal gas range, formula (15) is going to break down. For example P might start varying like ρ^2 , in which case the speed of sound will also depend on the pressure.

Then, for a given density, the speed of sound will depend on both temperature and pressure.

Q.: Coming back to the idea that the speed of sound is equal to the speed of molecules, I imagine that, outside the range of the ideal gas, that idea will also break down?

A.: It probably will. I never thought about it very much. What we studied is the depth that I have gone into it.

Let's move on to the harmonic oscillator. We shall first study a classical harmonic oscillator, and then a quantum oscillator.

Harmonic oscillator

Now we have not a free particle in a gas. We are going to put into the system – the heat bath, or whatever the system happens to be – a single harmonic oscillator. The harmonic oscillator will be part of the system. And we are going to treat the rest of the system as a heat bath. We shall look at the properties of that harmonic oscillator.

The harmonic oscillator under consideration could be a spring with a mass on it, figure 1, or anything else which naturally oscillates. It could be an electromagnetic wave in a cavity. What else oscillates? A crystal for instance. All sorts of oscillations take place in crystals. There are lots and lots of oscillating systems in nature. Almost anything will oscillate if you just disturb it a little bit.

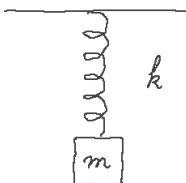


Figure 1: Harmonic oscillator.

Therefore when we study the harmonic oscillator, we are studying an extremely wide class of systems. We are studying anything which, when disturbed a little bit away from equilibrium, i.e. away from its ground state, will oscillate. And that covers almost anything that you could think about. So let's do the statistical mechanics of the harmonic

oscillator.

An interesting question is: What is the average energy of the oscillator? How does it depend on its mass? How does it depend on the spring constant? And so forth. Let's see if we can work out the answer to those questions.

So what do we do? We begin with a hamiltonian. That is, we begin with an expression for the energy. We are going to put it into the Boltzmann distribution. And then we are going to calculate average quantities for the oscillator.

The energy of the system is its kinetic energy plus its potential energy

$$E(\dot{x}, x) = \frac{m}{2}\dot{x}^2 + \frac{k}{2}x^2 \quad (16)$$

where x is the coordinate of the oscillator away from its equilibrium position, and \dot{x} the time derivative of x , that is the velocity of the little mass.

Caveat: k in equation (16) is not the Boltzmann constant, it is the Hooke's law spring constant. We are finished with Boltzmann constant now. It will fade in the background. We no longer will keep track of the Boltzmann constant.

If instead of velocity we use momentum, equation (16) becomes

$$E(p, x) = \frac{p^2}{2m} + \frac{kx^2}{2} \quad (16)$$

This is our expression for energy when the system in figure 1 is in state (p, x) . Now we are going to use it to calculate

the partition function Z for the system.

The partition function Z is our best friend in statistical mechanics. We use it, when we take its logarithm, to calculate all sorts of things.

In the discrete case, where the system is randomly in one of the states of $\Omega = \{\omega_1, \omega_2, \dots, \omega_i, \dots, \omega_n\}$, we saw in chapter 4 that

$$Z = \sum_i e^{-\beta E_i} \quad (17)$$

where E_i is the energy of the state ω_i , also called the i -th state. And β is the inverse of the temperature.

Take this opportunity to refresh your memory of what is the probability that the system be in the i -th state. You may also want to review how we defined or calculated α , β , the distribution formed by the P_i 's, S , the average energy E , Z , T , etc.

When the set of states Ω is a continuous phase space, each state ω is labeled by two coordinates x and p . Remember that x and p can each be of high dimensionality – although, in the present example, each is 1-dimensional. And the energy of one state ω is denoted $E(p, x)$. Then the partition function is

$$Z = \int e^{-\beta E(p,x)} dp dx \quad (18)$$

So let's write this partition function when the energy of the state (p, x) is given by equation (16):

$$Z = \int e^{-\beta \frac{p^2}{2m}} e^{-\beta \frac{kx^2}{2}} dp dx \quad (19)$$

As said, β is the inverse temperature. Now, for our system, what determines the temperature? The heat bath that it is immersed in. So if the oscillator in figure 1 was in the room where we are, it would be 300 degrees K. If it was in liquid helium, it would be 4 degrees. If it was in the middle of the sun... well, it would evaporate. But you get the point. T or β are determined by the heat bath.

When writing the partition function Z in equation (19), we readily used the fact that $E(p, x)$ is a sum of two terms, one depending only on p and the other depending only on x . This enables us to write the exponential beneath the integral sign as a product of two exponentials. Moreover, the integral in equation (19) factorizes into two integrals.

$$Z = \int e^{-\beta \frac{p^2}{2m}} dp \int e^{-\beta \frac{kx^2}{2}} dx \quad (20)$$

Now in previous situations the x -integration just gave us the volume of the container that the box was in, because there was no term in x in the energy. But here the x -integral is not just a volume.

Anyway, all we have to do now is to do these two integrals. Remember that multiplicative numerical constants in the partition function don't make any difference. Why not? Because when we take the logarithm of the partition function they correspond to additive constants. And since we are always going to be differentiating the log of the partition

function, we don't care about such numerical constants.

Let's look at the p -integral first

$$\int e^{-\beta \frac{p^2}{2m}} dp \quad (21)$$

We have already done it. But let's review the procedure to calculate it. We change variable

$$\beta \frac{p^2}{2m} = q^2 \quad (22)$$

Therefore

$$p = \sqrt{\frac{2m}{\beta}} q \quad (23)$$

The p integral becomes

$$\sqrt{\frac{2m}{\beta}} \int e^{-q^2} dq \quad (24)$$

The integral is a definite integral. Its bounds of integration are $-\infty$ to $+\infty$. It is a number. If you happen to remember, the number is the square root of π^3 . But anyway it is just a multiplicative numerical constant that we won't care about in equation (20).

Having taken care of the p integration, equation (20) becomes

³We proved it in chapter 4, footnote 9, using the trick of looking at its square and expressing the double integral in polar coordinates.

$$Z = \sqrt{\frac{2\pi m}{\beta}} \int e^{-\beta \frac{kx^2}{2}} dx \quad (25)$$

Now we play exactly the same game on the x integration. Let's change variable

$$\beta \frac{kx^2}{2} = y^2 \quad (26)$$

That means that

$$x = \sqrt{\frac{2}{\beta k}} y \quad (27)$$

Again, the y integral yields

$$\sqrt{\frac{2\pi}{\beta k}} \quad (28)$$

Finally the partition function is

$$\begin{aligned} Z &= \sqrt{\frac{2\pi m}{\beta}} \sqrt{\frac{2\pi}{\beta k}} \\ &= 2\pi \sqrt{\frac{m}{k}} \frac{1}{\beta} \end{aligned} \quad (29)$$

Remember that, for a spring, the quantity $\sqrt{k/m}$ has an important physical meaning: it is its frequency of oscillation, usually denoted ω . So we can write

$$Z = \frac{2\pi}{\omega} \frac{1}{\beta} \quad (30)$$

Of course the multiplicative constants in front of this expression won't matter, but at any rate that is the full expression of the partition function $Z(\beta)$.

Now let's calculate the energy of the oscillator⁴. How do we do that? We calculate the logarithm of Z and differentiate it with respect to β . That is the formula that we learned to use and to trust, see formula (19) of chapter 4.

$$\log Z = c - \log \beta \quad (31)$$

where the constant c is the logarithm of the coefficient in front of $1/\beta$ in formula (30).

Now what we want to do is we want to calculate minus the derivative of $\log Z$ with respect to β . We are cruising now. We are not thinking, just chasing symbols. We get

$$E = -\frac{\partial \log Z}{\partial \beta} = \frac{1}{\beta} \quad (32)$$

And of course this is the temperature. So we end up with

$$E = T \quad (33)$$

This is not so different than $E_{part} = 3/2 T$ of a particle in a gas.

The first question is: Why don't we have a factor 3 in formula (33)? That is because we are considering a one-dimensional oscillator, oscillating along a single axis. So

⁴By that we mean its *average* energy E .

there is no 3. Secondly, what happened to the two downstairs? Why isn't formula (33) $E = 1/2 T$? The reason is because there were two integrals, each of which, in the calculation of Z , gave a multiplicative factor proportional to square root of $1/\beta$. In the end, in the calculation of E , this produced the addition of two $1/2\beta$.

If, for whatever reason – more dimensions, more degrees of freedom –, there had been more Gaussian integrals in the calculation of Z , like the two integrals we had in formula (20), there would have been more terms $1/2\beta$ added in the calculation of the energy E .

What are the two kinds of energies, incidentally, that we are talking about? We are talking about the energy of an oscillator. It is kinetic energy and potential energy.

For an oscillator, the average kinetic energy and the average potential energy are in fact equal. And each one is equal to $1/2$ temperature. So the kinetic energy is $T/2$. The potential energy is $T/2$. Together they give us the temperature.

There are two interesting things to notice about formula (33). The first is that E doesn't depend on the mass of the oscillator. That was also true of the particle in the gas. At a given temperature, the kinetic energy of a particle does not depend on the mass of the particle.

The velocity of the particle did depend on the mass. The reason is that, at a given temperature, $1/2 mv^2$ is the same for all particles. So the higher m , the smaller v^2 .

The same is true in the case of the oscillator. The average

kinetic energy is one half the temperature. And it is independent of the mass of the oscillator.

The second observation is that E is also independent of the spring constant. As we saw, the potential energy is also one half the temperature. And k doesn't enter the formula.

That is a little surprising if you think about it. Indeed, imagine making the spring constant so large that no known force on Earth, or elsewhere, can stretch the spring away from its equilibrium point. One would think, under such circumstances, that it is not a spring, it is just a constrained object, which has a fixed length – the length possibly being zero – which cannot be changed.

In that case, any sensible person would say you can't excite it. You can't give it any energy. There is no way to start it vibrating. There is no way to give it any structure at all. You can't give it any kinetic energy because it is absolutely locked in place. You can't give it any potential energy because you can't pull it away from its rest point, no matter how much force you exert on it⁵.

Yet the formula doesn't seem to care. There always seems

⁵It seems that the reasoning consisting to pass to the limit, to try to show a contradiction, could also be applied to m alone: if m is too big, v would be zero, and there would be no kinetic energy. But in that case we would only have discovered that 0 times ∞ leads to unreliable conclusions.

However, the reasoning is not exactly the same on m and k . A huge m would mean a mass bigger than the Earth. Whereas a huge k is simply a perfectly rigid spring, which cannot be discarded as a ridiculous hypothesis like a huge mass. It is an ordinary laboratory equipment, at least from a classical mechanics point of view.

to be energy $E = T$ for every oscillator no matter how unphysical or how hard it is to get that oscillator to move.

So there must be something wrong in the analysis. The thing wrong is that we have ignored one very crucial feature of nature. We have not ignored anything about classical mechanics. These are correct classical mechanical conclusions. On the other hand it is crazy.

This occupied people around the turn of the XXth century. They were very confused by it. And the reason they were confused is because the ideal gas law and the energy per particle seemed to work pretty well when the particles were treated as a point particles.

Suppose, however, that the molecules could not be modeled as simple point particles. Indeed molecules can be for instance diatomic. Diatomic molecules can vibrate. They are little oscillators. In truth, they are rather stiff oscillators. They have a good solid spring constant. You have to give them quite a knock to start them vibrating.

Anyhow, the analysis seems to say that for a gas of diatomic molecules we ought to have not $E_{particle} = 3/2 T$, it should rather be perhaps $5/2 T$ to account for the extra energy due to the molecule vibrating like a harmonic oscillator.

But it was known that such extra energy wasn't there. So you could think of two conclusions.

1. One conclusion is that the diatomic molecule is not really a diatomic molecule at all. It is a point particle. But that is ridiculous.

2. The other conclusion is that we are doing something wrong or missing some ingredients of physics in our reasoning, which tends to keep that oscillator from having that much energy.

The missing ingredient is of course *quantum mechanics*.

Quantum harmonic oscillator

So what we want to do is the corresponding quantum mechanical calculation for a quantum mechanical oscillator in equilibrium, and compare the results with equation (33), which we obtained doing classical mechanics calculations.

How much do we have to know about quantum mechanics? Very little. We need the expression for the energy of a harmonic oscillator. The reader probably knows that harmonic oscillator energies are quantized. They come in discrete multiples of Planck's constant times the frequency. So the possible energy levels of the quantum oscillator oscillating with frequency ω are

$$n\bar{h}\omega \tag{34}$$

where \bar{h} is Planck's constant⁶. Equation (34) is a simplified version of equation (10.30) of chapter 10 of volume 2, in the collection *The Theoretical Minimum*, on quantum mechanics.

⁶ \bar{h} is the version divided by 2π , that is $\bar{h} = h/2\pi$.

The set of energy levels of a quantum harmonic oscillator can be represented as in figure 2. It is just a discrete set of states, all equally spaced.

That is all that we have to know about the quantum harmonic oscillator: the energy of a single harmonic oscillator comes in discrete integer multiples of $\hbar\omega$.

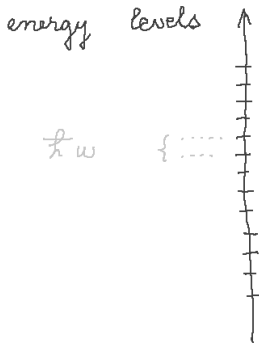


Figure 2: Energy levels of a quantum harmonic oscillator.

Now we can calculate the partition function for the quantum mechanical oscillator. As always, we sum $e^{-\beta E_i}$ over all energy levels E_i . For the classical oscillator, the space of states was a continuous phase space, so we replaced the sum by an integral, equation (18). But for the quantum oscillator, it is truly a sum.

$$Z = \sum_n e^{-\beta n \hbar \omega} \quad (35)$$

At first sight this looks complicated. But in fact it the simple series of powers

$$Z = \sum_n \left(e^{-\beta\hbar\omega} \right)^n \quad (36)$$

The first term, when $n = 0$, is 1. And the series has the form

$$1 + x + x^2 + x^3 + \dots \quad (37)$$

It is a geometric series. When $|x| < 1$, it converges and is simply equal to

$$\frac{1}{1 - x} \quad (38)$$

Since the exponent $-\beta\hbar\omega$ is negative, the basic exponential term in the power series expressing Z is less than one. The power series converges. And we can rewrite equation (36) as

$$Z = \frac{1}{1 - e^{-\beta\hbar\omega}} \quad (39)$$

This is the partition function for the quantum harmonic oscillator. There is a number of ways to write it. But this is fine.

What do we want to do with this partition function? We want to calculate the average energy. The general formula is

$$E = -\frac{\partial \log Z}{\partial \beta} \quad (40)$$

But here it is simpler to use the equivalent formula

$$E = -\frac{1}{Z} \frac{\partial Z}{\partial \beta} \quad (41)$$

What is the derivative of Z with respect to β ? There is an ugly thing in the denominator of the expression of Z , equation (39). When you differentiate a function which is in the denominator, you get first of all minus one over the denominator squared. And that is multiplied by the derivative of the function in the denominator. So

$$\begin{aligned} \frac{\partial Z}{\partial \beta} &= -\frac{1}{(1 - e^{-\beta \bar{h}\omega})^2} \frac{\partial}{\partial \beta} (1 - e^{-\beta \bar{h}\omega}) \\ &= -\frac{\bar{h}\omega e^{-\beta \bar{h}\omega}}{(1 - e^{-\beta \bar{h}\omega})^2} \end{aligned} \quad (42)$$

Now we have to divide by Z . That means multiply by $(1 - e^{-\beta \bar{h}\omega})$. It reduces the power in the denominator of equation (42).

$$\frac{1}{Z} \frac{\partial Z}{\partial \beta} = -\frac{\bar{h}\omega e^{-\beta \bar{h}\omega}}{1 - e^{-\beta \bar{h}\omega}} \quad (43)$$

This is minus the energy. We want plus the energy. So we change sign, and finally we reach

$$E = \frac{\bar{h}\omega e^{-\beta \bar{h}\omega}}{1 - e^{-\beta \bar{h}\omega}} \quad (44)$$

After all these calculations, by sheer luck we land on a positive number :-). Formula (44) is the answer for the average energy of the quantum harmonic oscillator.

It is to be compared with the average energy of the classical harmonic oscillator, which was simply $1/\beta$.

So what is the relationship between them?

Let's first look at high temperatures. High temperature is a situation where the classical theory is good. Each one of the little quanta of energy $\hbar\omega$ are thought of as very small. A classical spring has a great many units of energy. So in that sense the classical system in quantum units has a lot of energy. That must mean that it has very high temperature in some quantum sense.

Classical systems are quantum systems become classical when the temperature is very high. The quantization of energy then becomes unimportant.

Let's look at equation (44) at high temperature, that means when β tends to 0. Then the exponential term $e^{-\beta\hbar\omega}$ tends to 1. We can replace $e^{-\beta\hbar\omega}$ by 1 in the numerator. But of course we cannot do the same in the denominator. So we will use the second term in the Taylor expansion of $e^{-\beta\hbar\omega}$. Remember that

$$e^x = 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \dots \quad (45)$$

And the power series converges whatever x is, real or complex, with no bound on its magnitude. Anyway, when x is small we can take

$$1 - e^x \approx -x \quad (46)$$

So the energy given by equation (44) can be approximated as

$$E = \frac{\bar{h}\omega}{\beta\bar{h}\omega} \quad (47)$$

The next terms after $\beta\bar{h}\omega$ in the denominator of the right-hand side of equation (44) are proportional to β squared, β cubed, β to the fourth, and so forth. They are much smaller than $\beta\bar{h}\omega$ and they can be ignored when β gets very small.

So, at high temperature, equation (44) giving the energy becomes simply

$$E = \frac{1}{\beta} \quad (48)$$

At high temperature, it just reproduces classical physics: $E = T$, which is the same as equation (33).

Let's now ask about the low temperatures, that means when β becomes very large. Then in equation (44), the exponential term $e^{-\beta\bar{h}\omega}$ becomes very small. The denominator can be approximated by 1. And we can write

$$E = \bar{h}\omega e^{-\beta\bar{h}\omega} \quad (49)$$

We see that, far from getting just the temperature as in equation (48), we get something that is exponentially small when the temperature goes to zero.

We discover that quantum mechanics tends to suppress the energy of an oscillator when the temperature is low enough. At very low temperature, the oscillator doesn't behave like the classical oscillator at all. It has much less energy than

the corresponding classical oscillator would have had.

The question then is: Where is the crossover point? When temperature increases, where does it go from being quantum to classical? If we look at equation (49), we know that the crossover is where the exponential term $e^{-\beta\hbar\omega}$ goes from being small to being big.

The biggest $e^{-\beta\hbar\omega}$ ever gets is 1. But we can consider that the crossover between the low temperature behaviour and the high temperature behaviour is when the exponent in $e^{-\beta\hbar\omega}$ is about 1. So

- when $\beta\hbar\omega > 1$, quantum,
- when $\beta\hbar\omega < 1$, classical.

Re-expressed in term of temperature, this is: the behavior is quantum when $T < \hbar\omega$, and it is classical when $T > \hbar\omega$. The crossover point is when

$$T = \hbar\omega \tag{50}$$

If we step back and look at this result, it makes a lot of sense. It says that the crossover is when the energy of the oscillator, measured classically, that is with T , is equal to one quantum worth of quantum mechanical energy. If we try to make the temperature lower than that, then the oscillator has less than one quantum worth of energy.

The oscillator somehow doesn't want to have less than one quantum worth of energy. So when the temperature, or equivalently the energy, crosses this threshold downward, the expression for the energy, given then by equation (49),

becomes exponentially small.

To get a feel for the behavior of the energy given by equation (44), figure 3 represents the curve of E , for a given $\hbar\omega$, as a function of T .

When T is large, E is simply approximatley equal to T . And when T is small, we see how E gets exponentially closer to zero than T .

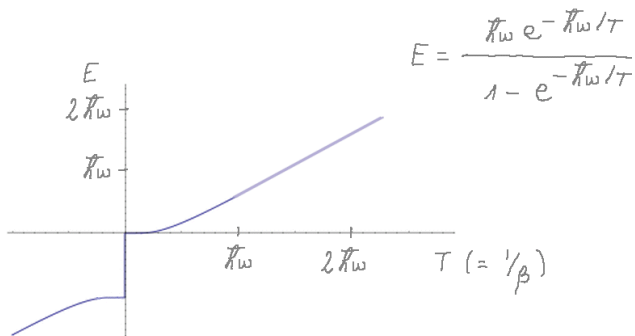


Figure 3: Curve of the energy given by equation (44) as a function of T . (Only the positive part is physically meaningful.)

Now we naturally know the answer to the puzzle about the diatomic molecule. When the temperature is very low we are deep in the quantum regime. Then the oscillator has exponentially small energy compared to what it would be classically.

On the other hand, when the temperature goes up, at some point the temperature will get larger than $\hbar\omega$. At that point we begin to activate the oscillator.

We can't activate the quantum oscillator unless the temperature is such that E corresponds to more than one quantum's worth of energy. That is what the threshold equation (50) says. At low temperatures, up to some temperature, the diatomic molecule behaves like a monoatomic molecule. The molecule doesn't vibrate like an oscillator.

When we raise the temperature, past a threshold where the temperature becomes bigger than the quantum of energy, the molecules start to behave like genuine diatomic molecules.

Let's turn to the characteristics of the molecular oscillator. Its stiffness determines ω . Remember that $\omega = \sqrt{k/m}$. So the stiffer the oscillator, the higher the frequency of oscillation. Therefore the stiffer the oscillator, the higher the temperature of the crossover point $\hbar\omega$ will be.

Molecules are pretty stiff on the scale of rubber bands. So the threshold temperature is fairly high.

In fact, before we cross the threshold, something else happens. Molecules start to rotate. Other things start to happen. Our calculations are just an illustration of happens when we go from quantum to classical behavior.

Then what happens if we raise the temperature even more? We activate even worse things. The atoms themselves start vibrating. Eventually they will get ionized. When an atom becomes ionised, it doesn't behave like an atom at all. If it is an hydrogen atom, it start to behave like an electron and a proton.

This is an example of quantum mechanics solving a deep problem that physicists have been very confused about around the turn of the XXth century.

A similar problem came up in the following context. Solids are simply systems of molecules where the molecules are arranged in a lattice. They vibrate around their average lattice positions. So, beside its possible internal structure deformation – which usually is much stiffer –, each molecule is also sort of an oscillator around its lattice position.

The energy of these oscillators as a function of the temperature determines the specific heat, or heat capacity, of the crystal. However the theoretical estimations of their values by classical physics, for various crystals, were much larger than the actual measurements.

It was Einstein who figured out what was going on. He explained that the quantum statistical mechanics of oscillators essentially suppressed them until the temperature got high enough.

Einstein published his paper on the specific heat of solids in 1907. He had already thrown in some of the seminal ideas that would lead to quantum mechanics, for instance his explanation, in 1905, of the photoelectric effect with photons. But his paper on specific heat came twenty years before people understood the foundations of quantum mechanics. In 1907 Einstein just knew that the quantum energy levels had to be proportional to frequency and to n .

So it is another of Einstein's contributions – one of his minor ones.

Before we leave oscillators, let's step back to look at what we did and found.

We went from the classical harmonic oscillator to the quantum harmonic oscillator, because we followed the chronology of scientific progress.

Staying in classical physics, we saw that we could explain pretty well the behavior of an oscillator plunged in a heat bath. But the theory ran into trouble for very stiff oscillators. It predicted too high energy at low temperatures.

Then we went into quantum mechanics, staying at a very simple level, which corresponded to what was known in the first decade of last century. No fancy complex vector space describing the states of a system, no Hilbert space, no distinction between states and observations, no fundamental randomness in the measurements made on a given superposition of states, when the state is expressed in the basis attached to the Hermitian operator of an observable, etc. See the volume 2 in the collection *The Theoretical Minimum*, on quantum mechanics. We just assumed the quantification of the levels of energy. And we saw it was sufficient to solve the problem.

As we already pointed out in volume 2, the right way to go – or at least the logical way to go – is to study first the quantum mechanical description of nature, and then go to the classical description as a limit case when systems are

made of a large number of particles, and the temperature is not too low.

The classical physicists were right in using a continuous phase space, calculating the partition function as the integral in equation (20), and showing that $E = T$. They just had a problem at low temperature, that is low energy. The oscillator had too much energy compared to experimental observation.

Elementary quantum mechanics solved the problem for temperatures below the threshold of $\hbar\omega$ – which happens to be high, and that is why classical physicists were aware of the problem and terribly puzzled by it. But above the threshold, quantum mechanics confirms that $E = T$. Then it does not depend on the stiffness of the oscillator or anything, only on the temperature.

In fact, elementary quantum mechanics did not change much the partition function. In the classical case we obtained equation (30), reproduced below

$$Z = \frac{2\pi}{\omega} \frac{1}{\beta} \quad (51)$$

In the quantum case equation (39) gives

$$Z \approx \frac{1}{\hbar\omega} \frac{1}{\beta} \quad (52)$$

If we remember that $\hbar = h/2\pi$, we see that quantum mechanics simply replaces 1 by h in the denominator of the classical partition function. And since multiplicative constants in partition functions don't play any role, at first sight they

are quite similar.

The difference appears when we calculate the energy, with equation (30) in the classical continuous case, leading us to $E = T$ no matter what; and with the equations (39) to (43) in the quantum case, leading us to the more complicated equation (44).

Once it is above a quantum threshold, the energy is just equal to the temperature and it does not matter what the spring constant is.

Question: What happens to the oscillator in terms of kinetic and potential energy once we are above the threshold temperature? Answer: It speeds up and it makes larger and larger oscillations as we increase the temperature.

The point is that, as we heat it from low temperatures to high temperatures, the temperature at which the classical behavior sets in, at which the energy paradoxically becomes independent of how stiff the spring is, and becomes simply $E = T$, gets to be a higher and higher temperature as a function of the spring constant.

So if we take a very very stiff spring constant, let's say $k = 10^{10}$, then, yes, at sufficiently high temperature, it will have an energy equal to its temperature. But at any kind of ordinary temperature, it would essentially have no internal energy. It would be the motionless and energyless solid we expect it to be – leaving aside the motion due to the heat bath.

Look again at figure 3, to see that it is exactly what the curve of E shows. The threshold for the oscillator to start behaving classically would get higher and higher as the spring constant got stiffer.

This leads us to an interesting observation.

Degrees of freedom

The transition from quantum to classical behavior, when we cross a threshold in temperature, is interesting because it has the following consequence:

As we heat a system, we start to discover that it has more and more degrees of freedom.

To understand this better, let's make a crazy molecule. It consists of a spring with two masses connected to it, at each of its extremities. The spring is rather stiff. Furthermore the mass on the right is itself actually made of two little masses linked by a secondary spring, much stiffer than the main one, as shown in figure 4.




Figure 4: Molecule whose number of degrees of freedom increases with temperature. The secondary spring on the right is much stiffer than the main spring in the middle.

How does it behave at low temperatures? At low temperatures there is not enough temperature to overcome the quantum threshold for either of the two oscillators to start oscillating. So the whole molecule behaves like a point mass molecule. It behaves with the energy $3/2 T$, like any other point mass particle in the heat bath, see equation (64) of chapter 4.

Then we heat it up to the point where we activate the main oscillator in the middle. But still we don't give it enough temperature to activate the secondary super duper strong spring on the right. So the molecule simply behaves like a diatomic molecule, with for instance $5/2 T$.

We heat it more. Eventually we reach the threshold where we also activate the secondary spring. The molecule starts to behave like three masses connected by springs. Instead of the $5/2 T$ of the diatomic molecule, it behaves now with for instance $7/2 T$, or whatever the right number is.

The lesson to draw from this example is that, as we heat the system, more and more degrees of freedom become activated. And we start to discover the complexity of things.

We don't discover the complexity of things at low temperature. When the temperature is low, we discover the quantum behavior of things.

As we heat it, we start to discover the complexity of the system under study. Degree of freedom by degree of freedom, unfreezing themselves from the quantum constraint, the system starts behaving more and more classically.

The reader may realize that this has a lot to do with the blackbody radiation. We will come to that. But we won't relate it right now to the blackbody radiation.

For the moment, let's relate it to something else: a violin string.

A violin string consists mathematically of a large number of oscillators. The oscillators are the string oscillating in its various harmonics, figure 5.

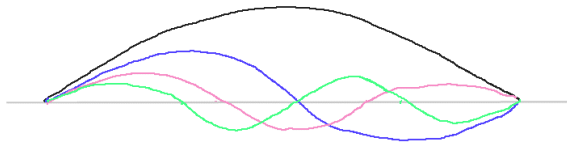


Figure 5: The first four harmonics of a violin string vibrating.

In other words, from a fundamental mathematics point of view, a vibrating string is basically an infinite number of harmonic oscillators, of shorter and shorter wavelengths.

If each one of these oscillators had energy equal to the temperature, as given by equation (33), and there were an infinite number of them, then that string would have an infinite amount of energy when it came to thermal equilibrium.

So something is wrong.

The answer to the riddle is that most of these oscillations have very high frequency. Indeed the shorter the wavelength the higher the frequency of the oscillator. And therefore most of these oscillators, at any given temperature, are frozen out by quantum mechanics. For those frequencies, $T < \hbar\omega$. So they are not free to oscillate.

As the temperature goes up, more and more of these harmonic oscillations appear, representing more energy of the string.

That is another paradox similar the one discussed in the previous section.

Questions / answers session (2)

Q.: If we want to know the values of these numbers k , T , $\hbar\omega$, etc., in the transition from quantum to classical behavior, can we observe and measure them experimentally,

somewhere, somehow?

A.: Yes, absolutely. They can be observed and measured experimentally. It is done just studying the specific heat of crystals. The experimental phenomena are obvious. And once physicists understood what was going on and had the right model, the calculations were easy.

It was obvious to experimental physicists around 1900 that there was something wrong, because their crystals did not behave according to classical mechanics.

So yes it is extremely easy to see the transition, and to make measurements about it.

Q.: If we heat the crystal, we can make it change its behavior?

A.: Yes. If you heat the crystal, the specific heat starts to behave as if there was a certain number of T , or more precisely a certain number of T per oscillator. This is really easy to see.

Q.: This would translate into the partition function? We could handle the multiple oscillators?

A.: Yes. The partition function would be the product of all the partition functions, one for each oscillator. The logarithm would be the sum of the logarithms.

The logarithm being the sum implies that the energy of the system is the sum of all the oscillators energies⁷.

But those whose frequency is too high, above the threshold temperature, those are in the quantum regime. And they are not appreciably excited.

Q.: In describing the quantum levels of energy of the quantum oscillator, equation (34) and figure 2, in order to compute the partition function, why did you ignore the ground level energy?

A.: You are right, we did ignore the ground level energy. Instead of using the proper equation

$$E_n = \bar{h}\omega (n + 1/2) \quad (53)$$

we simplified it into equation (34) reproduced below

$$E_n = n \bar{h}\omega \quad (54)$$

The reason is that it doesn't matter. Adding a constant to every energy levels is equivalent to multiplying the partition function by a constant. And the result is no change in anything of interest.

⁷When we calculate a partition function we take into account the exact energy level of each possible state of the system. Then, when we talk about calculating the energy from the partition function, we mean the average energy. So the average energy of the system, would be the sum of the average energy of each oscillators.

Let's go to the next topic, which is the second law of thermodynamics. We will first of all explain what is puzzling about the second law. Then, after a third questions and answers session, we will give one statement of the second law. After that we will discuss further trajectories in the phase space, chaos and recurrence.

Second law of thermodynamics

Why is the second law a puzzle? It is a puzzle because it says that the world is irreversible. In isolated systems, something called entropy always increases, or stays the same.

In fact you would have to be infinitely careful, if the isolated system you examine is big and not in equilibrium, to keep its entropy from increasing.

Remember that the entropy of a system, unlike its energy, temperature, pressure, etc., is not a characteristic of the system in itself. It is a characteristic of our incomplete knowledge of the exact state the system is in. According to equation (30) of chapter 1, the entropy of a system, when we know exactly in which state it is, is zero.

But in general, because it is difficult to follow exactly all the positions and motions of all the particles of a big system – for instance of a gas in a vase –, at time t the system has a certain entropy. Moreover the entropy of an isolated system will increase, unless it stays the same. For instance the entropy of a gas in a vase isolated from the exterior, and not yet internally in thermal equilibrium, will increase until it

reaches a maximum value, corresponding to the Boltzmann distribution over all its possible states, see chapters 3 and 4.

The puzzle stems from the fact that, on the other hand, Newton's equations of motion are completely reversible. So anything that can happen in one direction of time can happen in the opposite direction of time – implying, in particular, that the entropy of an isolated system should be able to decrease spontaneously.

So there seems to be a contradiction. At the least, we can say this:

There is a tension between reversibility of the fundamental equations, and irreversibility of the observational properties of complicated systems.

That took some time to sort out. And many people are still confused by that. So let's talk about it.

Let's imagine some phase space. A phase space is the space of states of a classical system. It has got a lot of dimensions. It has got a lot of particles, or a lot of whatever it is made of.

But, as usual, we are going to represent it simply as two-dimensional, figure 6.



Figure 6: Phase space of a system.

One dimension is the coordinates. The other dimension is the moments. There are a lot of them. So this is the sketchy representation of a high dimensional space.

We start out, at time t_0 , with some probability distribution over a subset of all the possible states. For example let's start out with some probability distribution over the subset represented in figure 7.

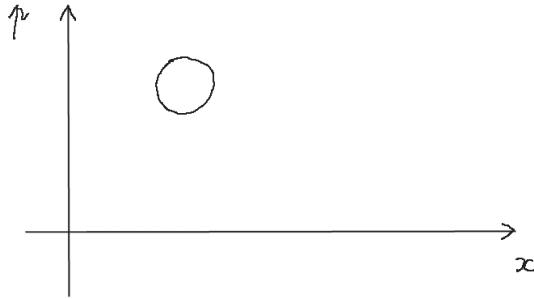


Figure 7: Initial subset of possible states.

The picture doesn't show explicitly the probability distribution, or density, which would be in a third dimension toward us for instance, but we can imagine it to be a uniform distribution. So there is zero probability for the system to be outside the blob, and uniform probability to be inside of the blob in figure 7.

This means that we know nothing about the exact state of the system, except that it starts inside the blob. Since we don't know anything more about it, the probability density is constant within the blob.

Do you remember what is the entropy of that system? Go back to chapter 1 if you need to refresh your memory. When the distribution is uniform, the entropy is the logarithm of the volume of the blob, equation (31) of chapter 1. More generally, it is $-\int P \log P$, equation (40) of chapter 1.

It is not yet the maximum of anything⁸. It is not the result of some maximization under constraint. We just happen to know that the system is within that blob. We did our best possible measurements and we discovered that it is there.

In particular, *the system is not in equilibrium*. We are not talking about the entropy in equilibrium⁹, we are talking about the entropy of a certain probability distribution.

The probability distribution is handed to us. Somebody told us that the momenta and the positions of 10^{23} particles correspond to a point within the blob¹⁰. But we are not told their exact values, we are only told they are in there, with a uniform distribution.

This situation is not that unrealistic. We might start with all the air molecules in some air tank in the corner of the room. Then we open up the tank, and let the air go.

Originally, at time t_0 , we knew that the molecules were all confined in a small volume in the corner. That corresponds to some section of the phase space. Moreover we had a pretty good idea of what their energy was, because

⁸Except, it is the maximum of our lack of knowledge within the blob.

⁹Notice that for a system in equilibrium the entropy doesn't change with time, although the molecules still zip around all over the place at a few hundred meters per second, and keep colliding with each other. It is a pandemonium, and it is known that it has temperature :-). So in thermodynamics, equilibrium doesn't have the same meaning as in Newtonian dynamics and statics.

¹⁰We use the number 10^{23} to mean "a lot", because, as the reader remembers, 6.02×10^{23} is the number of particles, under normal conditions of temperature and pressure, contained in a gas in vase of volume 22.4 liters

we knew what the temperature was in the tank. So we had a pretty good idea of what the momenta of all the particles were. That gives us a blob¹¹. There might not be exactly a uniform density of probability over the blob, but it would certainly correspond to some blob-like configuration.

Make sure you clearly see that the more we know, the smaller the volume of the corresponding region in the phase space is. Recall that, when the probability is uniform, the entropy is the logarithm of the volume. So, as the volume gets smaller and smaller, the entropy goes down.

Now, what happens to the blob as time goes on? Remember Liouville's theorem, which we studied in chapter 1, and earlier in volume 1. It says that the volume of the blob stays the same. Therefore the entropy stays the same.

It is just another way of saying that if we were able to know the exact state the system is in at some time t_0 and we were able to follow what it becomes, it would stay unique in the future. And that would be true also going backward in the past. Each state comes from a unique state, and goes into a unique state. And if we could follow the system we would see it.

The discrete version of Liouville's theorem is easy to figure out intuitively. If the system at time t_0 is, with respective probabilities $1/3$, $1/3$ and $1/3$, in the three possible states ω_1 , ω_2 and ω_3 , then, at a later time t_1 , it will still be in

¹¹Don't mix up the small volume in the corner of the room with the corresponding blob in the phase space. What would be the blob corresponding to all the molecules in the air tank? What would it be if we knew exactly the positions and momenta of all the molecules?

three possible states, not necessarily the same, but with the same equal probabilities $1/3$. And it is even true with non-uniform probabilities p_1 , p_2 and p_3 .

That is one meaning of the following law, which we called the *law number -1*, because it is even more fundamental than the laws number 0, 1, 2, etc. :

In nature information never gets lost or erased.

We observers may lose information, but Nature doesn't¹².

Liouville's theorem – saying that the collection, or blob, of possible states may change from one time to the next, but that the shape of the distribution doesn't – seems to imply that entropy doesn't change. And it is true. This is the correct view of what is called *microentropy*.

But there is another concept of entropy which does increase. This other concept of entropy has to do with something called *coarse graining*.

Typically our experiments are such that we cannot resolve¹³ all points in the phase space. Even if our experiments were ideal, i.e. perfectly precise, we still could not tell one state

¹²In this statement, like Animists, we personify Nature because it is a convenient way of saying. But the philosophy of the co-author is phenomenology: Nature is a bunch of raw phenomena, and we humans build models to represent them within coherent structures. Those even have some predictive power, although as many philosophers and scientists pointed out "to predict is not to explain" – remaining vague on what to explain is.

¹³To resolve, here, means to distinguish, like in the resolution of a photograph.

from another state if they were too close. There is some maximal resolution of the state that we can achieve.

In quantum mechanics, it turns out that a limit to the resolution is naturally built in. The smallest detail that we can know, in other words the smallest uncertainty, is given by Planck's constant. Remember that $\Delta X \Delta P \geq \hbar/2$, see chapter 8 of volume 2 in the collection *The Theoretical Minimum*, on quantum mechanics.

But in other situations it may just be that our resolving powers are not good enough.

So, instead of talking about points of phase space, we talk about *resolved points*, or *coarse grained points*, in phase space.

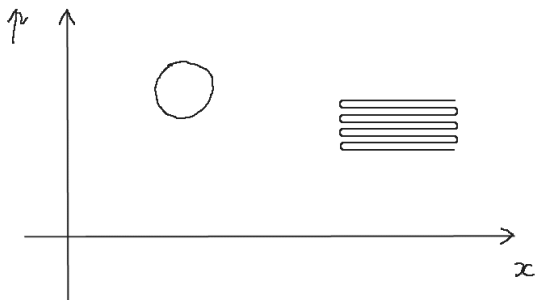


Figure 8: Example of dynamics of a blob: from the patch on the left it becomes a long folded snake on the right, *with the same volume*.

Let's make an example. Supposing we started with the blob in phase space shown in figure 7, or on the left part of figure

8, and it has a certain volume. Let's follow it.

We tell what it evolves into, not solving any mathematics problem; we just make up an example of dynamics which preserves the volume. So suppose the dynamics transform the initial blob into a long snake as shown on the right part of figure 8.

The snake is extremely thin, and the folds of the snake are extremely thin. Yet the volume of the snake is the same as the volume of the initial blob.

But unfortunately, because of our limited resolution powers, we cannot tell a phase point from a neighboring phase point when they are too close.

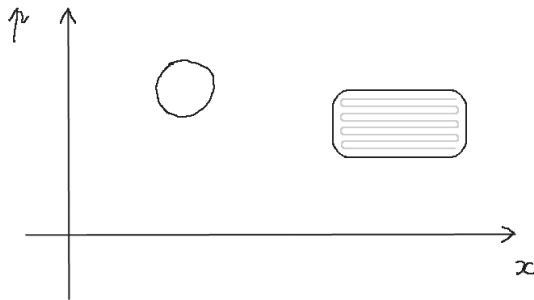


Figure 9: We view the snake as a blob.

What we do when this happens is we take the phase points of the snake and "blobify" them. That is, we replace them by small blobs. They all overlap. And the result is that

the initial blob now became not a snake but a bigger blob, figure 9.

We don't violate Liouville. The snake in figure 8 is still consistent with Liouville. We are just losing information in figure 9, because we can't keep track of the fine structure of the snake. We have a fuzzy resolution. It is as if, so to speak, we were looking at these things with somebody else's eyeglasses. So instead of the snake we see a new larger blob.

If we were asked "where in phase space is the system now?", we would answer that, as far as we can say, its position is uniformly distributed over the larger blob.

The crux of the matter is that we couldn't follow the system carefully enough. So we lost information.

If we could follow the system with infinite precision, and we could figure out the initial probability density precisely too, then we would know that the system evolved into a point somewhere on the snake. The volume would not change, representing the fact that we would not lose information.

But in real situations, we can't distinguish one state from a nearby state. So we do this process of coarse graining¹⁴, meaning by that, replacing points in phase space by blobs in phase space.

Once we do coarse graining, the volume increases. The snake, which was actually a blob with the same volume as

¹⁴Process may be too strong a word. It does require any action on our part, except recognizing that we lost information, and now the blob representing our lack of knowledge is larger.

the initial one, becomes a blob with a bigger volume.

Here is an image which can help figure out what is going on in the way we perceive the dynamics of the system. Imagine a piece of cotton wadding, figure 10.



Figure 10: Cotton wadding.

What is the volume of this piece of cotton wadding? If we mean the strict volume of the cotton fibers all added up, we can squeeze them down to something quite small.

If we let the cotton fibers grow, the strict volume of the cotton wad doesn't change very much incidentally. But the cotton wad sort of grows out. And if we look at it again, it has a bigger apparent volume¹⁵.

¹⁵In the case of a blob of phase space points, when it evolves its strict volume doesn't change at all – that is Liouville's theorem –, unlike the wadding when we let the fibers grow. But the image of the wadding is still appropriate to compare the strict volume of the snake to its apparent volume after coarse graining.

Returning to coarse graining, if we look at the cotton wadding through somebody else's glasses and we can't see the cotton fiber, we will say that the volume of that piece of fiber stuff has increased substantially.

In essence, that is what happens in the dynamics of the blob and the perception we have of it, figures 8 and 9.

Questions / answers session (3)

Q.: Are you saying then that classically it only appears to increase, whereas quantum mechanically it does genuinely increase because of the uncertainty principle?

A.: Let's forget quantum mechanics.

We are in classical mechanics. The phenomenon is just due to the limits of our resolution power.

Remember that entropy is not really a property of a system. It is a property of the system and what we know about the system.

So if we could follow in infinite detail every point in the phase space, and all we knew in the beginning was that the system was inside some blob, but we did not exactly where, and our lack of knowledge was represented by a uniform probability distribution, then we would say afterwards it is inside some other blob elsewhere with a comparable uniform probability distribution, and therefore the same volume.

We took the example where afterwards the new blob had the shape of a snake. So again, *if we could follow every point exactly*, we would know afterwards it is in the snake some-where. But we would not know exactly where it is in the snake. And entropy would be conserved.

But if we take into account that we can't tell one point from another, when they are too close to each other, because of limits of resolution, then we wind up saying: "Look I know much less about where it is that I did to begin with." Then entropy went up.

Q.: So classically we can imagine that we might be able to keep track of the system in such a way that the entropy would remain the same.

A.: Yes, that right. But the problem is that, given any degree of resolution – whatever our resolution power is –, we will run into the same difficulty.

Let's suppose it was good enough to resolve the snake after 10 seconds. So, for example, it is good enough to resolve the snake in figure 8.

But then, because of the way systems evolve in phase space, they start to grow even littler snakes, and littler snakes, and littler snakes...

Whatever our degree of resolution, this "snakiness" will appear. Let's give its right mathematical name to the phenomenon: it is called chaos.

Let's now wrap up our description of the second law in terms of increasing chaos and coarse graining leading to a loss of information.

Second law of thermodynamics (con'd)

The way systems evolve in phase space, from an initial blob, is they start to spread out sort of tentacles, without changing the volume. They grow fibers if you prefer. And the fibers grow thinner fibers. And those thinner fibers grow even thinner fibers – the total volume staying always the same.

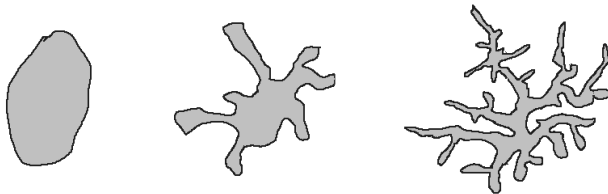


Figure 11: Chaotic growth, without changing the total volume.

Eventually, no matter how good our resolution is, we will start seeing that the entropy increases. And that is what the second law of thermodynamics says.

One of the numerous equivalent statements of the second law of thermodynamics is:

The entropy of a closed system not in equilibrium increases until it reaches a maximum. At that point the system is in thermal equilibrium. And the uncertainty of our knowledge then corresponds to the Boltzmann distribution over all the possible states of the system.

In the next chapter we will refine the above statement, adding that this evolution of the entropy is not absolutely sure but has an *extremely high probability*.

Let's continue investigating the interplay between the deterministic trajectory and our loss of information.

Trajectories in the phase space, determinism, chaos and recurrence

Let's imagine some phase space. We started out knowing a great deal. That is, we started out knowing that the system was somewhere inside some small blob. Then the system evolves in the phase space.

Let's first of all talk about how big the entire available region of phase space can be. Suppose the system is a gas made of a large number of particles in a box or a vase.

That means the x 's are bounded. For instance, in a one-dimensional analog, we would say that the available phase

space is bounded within an interval.

Can the momenta be anything? No. Usually the energy also constrains how much the maximum of the p 's can be. Even if all of the energy of the system went into one particle, that one particle would still itself have a finite energy.

Therefore the effective region of phase space which is available to us is some big box, figure 12.

Maybe the shape of the region is not a high dimensional rectangle. It has got some shape. But it doesn't extend to infinity in any direction. It is finite.

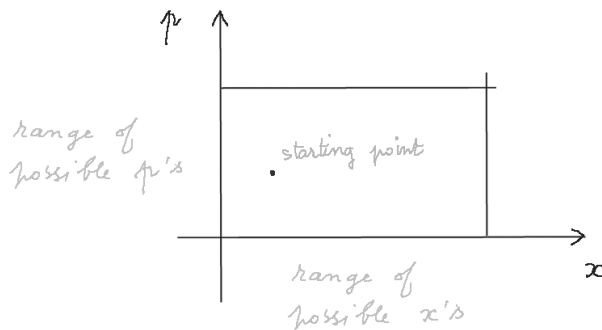


Figure 12: Maximum possible region in the phase space, for a gas in a box or a vase.

Now the phase point starts someplace. The exact precise point is shown in figure 12. And the system starts to evolve.

For most systems, what happens is that the subsequent path is *extremely sensitive* to the starting point. So if we started initially just a little bit away from the first point,

the paths would look the same for a little while, and then would be quite different, figure 13.

That is the essence of mathematical chaos: an equation of motion is classified as chaotic – or a system is classified as chaotic – if a small difference in the initial conditions leads quickly to quite different trajectories of evolution.

An example would be billiard balls. Instead of the gas in a vase, think of it in two dimensions. The vase is now an idealized billiard table with no friction. The walls are the cushions. And the molecules are the billiard balls¹⁶.

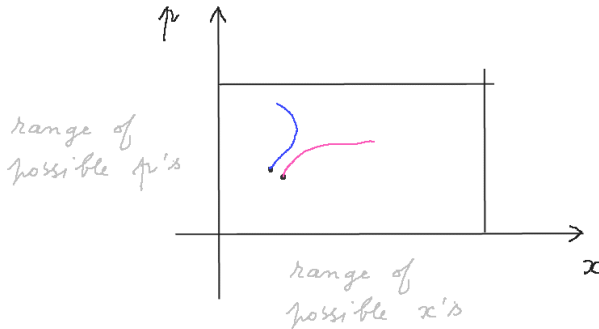


Figure 13: Sensitivity to initial conditions.

How many readers play billiard? Do you play pocket pool? Chicago? Straight pool? Let's leave three-cushion aside...

¹⁶Again, do not mix up the billiard table and figure 13. When we represent a billiard table viewed from above, we represent only x_1 and x_2 . We don't represent the p 's. Whereas figure 13 represent all the x 's as the horizontal axis, and all the momenta as the vertical axis. So a trajectory on the billiard table doesn't have the same meaning as a trajectory in figure 13.

too hard, too much chaos :-)

So we arrange the balls in a certain fashion and start with a cue ball. If, on two different occasions, we reproduce the same initial configuration extremely accurately and we shoot exactly the same, then we know that exactly the same things will happen. That is indeed what Newton's laws of physics say.

However, everything else being kept perfectly fixed, what if we are a tiny bit off in our aim? Well, if we are really just off by a truly tiny amount, the difference in the trajectory of the whole system will be small for a while¹⁷. But sooner or later it is going to start to diverge exponentially, because errors compound. Sooner or later, the two trajectories start-ing from nearby points in figure 13 are going to deviate from each other.

As we said, that is the character of what is called *chaotic systems*. And most systems in nature are chaotic.

So that means that, no matter how close we start together, we will eventually depart.

Moreover a small initial blob will grow into into a lot of phase space, even though its volume in the phase space strictly speaking stays the same, figure 11. The phase space subset which the intial blob evolves into gets extremely complicated and ramified.

Eventually, if we wait long enough till the system gets to

¹⁷Be it the spatial trajectories of the balls on the table bed, or the more abstract trajectory of the system in the phase space.

thermal equilibrium, the cottony stuff basically fills up the *entire possible region* of the phase space, figure 13. And if we do any coarse graining at all, it covers it with uniform probability.

Now let's start from any point in the phase space. To be practical we can start for instance from the usual initial configuration of the fifteen colored balls arranged at rest in a triangle and the cue ball with an initial impulsion. But any other point where the sixteen balls have a position and a velocity would do. Let's follow the trajectory of the system, figure 14.

Supposing the billiard is ideal, the table bed has no friction, for simplicity there are no pockets, and the balls run forever, then if we wait long enough we will come as close as we want to any possible configuration of positions and velocities of the sixteen balls in the phase space. Every possible configuration in the phase space will get visited – not exactly but very close.

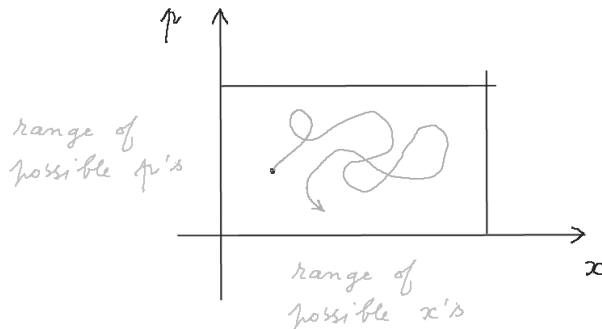


Figure 14: From any starting point, if we wait long enough the subsequent trajectory will get as close as we want to any point.

That applies to any starting point, even those special¹⁸. For instance, going back to the example of the gas in the room, if we start with all the molecules in the corner of the room with a certain definite velocity going off in some direction, they will follow a trajectory which will fill up, in the sense explained, the whole possible phase space.

So, after a while, we don't know anything anymore. The system has reached its maximum entropy.

Nevertheless, in a completely unpredictable and random way that we can't foresee – because we can't follow the system with enough detail –, the trajectory of the phase point will eventually come back as close as we like to the initial phase point where all the molecules were in the corner.

In the billiard analogy, if we start with the triangle of colored balls at rest and the cue ball with an impulsion, the balls will scamper in all directions, they will run all over the table bed, bumping into each other and against the cushions, but if we wait long enough at some time in the future, they will pass through being all arranged as initially and with only the cue ball having some velocity.

How long will it take? Well it is not infinite. It does depend on how precisely we require to reproduce the original configuration. If we required infinite precision, for most initial points it would never come back. But if we allow a tolerance, even very small, then the time that it takes to pass back through that phase point to within tolerance

¹⁸Notice the special meaning of special. What does special mean? In the theory of complexity, it means something describable simply. For instance: all the molecules in one corner.

epsilon is finite.

So, on the one hand, the entropy increases. It increases *in the coarse graining sense* because we lose track of the exact position and motion of every particles. On the other hand, that does not say that we can't return to the original configuration.

It is possible to estimate how long it will take for a system of many degrees of freedom to come back to within a certain tolerance of any phase point. Take the air in the room, put it all up in the corner, with its kinetic energy, all the molecules for instance going in the same direction with the same speed, that is, a very special state. Let them go. They shoot off to the other wall, bounce off that wall, bounce off other walls, start hitting each other, etc. They go zipping all over the place.

Such a system is very chaotic. So, after a few minutes at the most, they fill the room in a normal pattern of thermal equilibrium¹⁹. We observers sit and wait, and wait, and wait. How long will it take to see all the molecules back again in one corner?

¹⁹To get a feel for what thermal equilibrium looks like: In the air in the room, under normal conditions, there are thirty thousand billion billion molecules per liter. At thermal equilibrium, each molecule runs freely, at five hundred meters per second, over a few nanometers before hitting another molecule, or the walls, or our skin.

Despite this pandemonium like picture, this gives us a nice feeling of cozy temperature. Our body being at 37° Celsius, we contribute however to heating the room. We are in a constant dynamic equilibrium of transferring heat generated by our basal metabolism, which we need to refuel once in a while with a fresh chocolate brownie.

We can be certain that, if the experiment lasts long enough, and the room is a closed system, or at least is kept at constant temperature and energy, the molecules at some point will come back to their initial state. The question of how long it will take can be answered.

The answer is: a long time:-)

Questions / answers session (4)

Q.: Don't we lose information because it is not a reversible process?

A.: Your question mixes up two things.

On the one hand we do lose information, because of the limitation of our resolution power. If you can look at the trajectory of a minuscule ladybug on the table with perfect eyesight, you always know exactly where it is. But then if you put on the spectacles of somebody else which blur badly your sight, you will no longer know exactly where is the bug. You will only perceive a blurry red blotch. And you will know the bug is somewhere in the blotch. The same happens to us in the phase space of positions and momenta.

On the other hand the process is perfectly reversible. So that is not why we lose information. The process is completely reversible: anything that can happen, can also happen with time flowing in the opposite direction.

But we will talk more about it. We will talk more about the second law. In chapter 8, we will talk about reversibility, and how this paradox gets resolved.

It was one of the great classic paradoxes in the history of physics which, as you undoubtedly know, puzzled a lot of people. It drove Boltzmann to distraction. I don't know if that is what drove him to suicide, but it certainly depressed him.

In fact Boltzmann eventually came to the right answer: it is not that entropy always increases, but, given any configuration, it is most probable that the next configuration will have more entropy. In other words, entropy probably increases.

The probability for the configuration aiming itself back to some unusual, unexpected configuration, that is always small.

We will give more examples, and try to make this phenomenon a little more intuitive in the next lesson. We will go into the second law in more detail, and at least provide some kind of explanation.

Q.: When you say that most systems in nature are chaotic, is that a principle or an empirical fact?

A.: Well there is an infinite number of chaotic systems. And there is an infinite number of non-chaotic systems.

Now we have to start talking about measures of the chaotic character of a system, measures on hamiltonians and things like that.

It is a good question. I'm sure that the mathematical people have come up with a correct measure of what is a chaotic system vs a non-chaotic system.

Not all systems are chaotic, that is for sure. In many cases you can perturb them a little bit, you can change them a little bit, and they will still remain non-chaotic. It is not the sort of situation where, if you have a system and you change the initial conditions or the rule, meaning the hamiltonian, a very tiny bit, it will evolve quite differently.

What does chaotic mean? It means that the phase trajectories tend to separate and become quite different, like in figure 13, after a certain amount of time. I didn't want to define it precisely in this lesson, but it is this phenomenon of close by points in phase space following each other for a while then exponentially departing. So that the system becomes effectively unpredictable.

It is a situation in which predictability effectively breaks down, because, in order to predict a system for a length of time t , your precision with the initial conditions and your knowledge of the rule of evolution of the system, has to get better and better as you increase the time over which you want predictability. And necessarily at some point this exceeds the experimental precision we have access to²⁰

²⁰In our mathematical and physical models, we often – most of the time – use precise measurements, precise positions, exact numbers expressing one thing or another. This square has a diagonal of $\sqrt{2}$. The

The point not to oversee is that the system is and remains deterministic.

Let's take the weather for instance. Just define it as a collection of molecules doing whatever they are doing instantaneously. And let's ask how long can we predict the weather? It depends on how precisely we know the initial conditions (ideally, the exact position and velocity of every molecule!).

The longer the period of time we want, the more precise the initial data must be. Even longer, even more precise.

Now that is quite different than the harmonic oscillator.

energy required to lift this mass over one meter is 1 Joule, etc. We are so accustomed to doing that, that we don't see what could be wrong with it. There is nothing wrong with it, but *it is just a model*. Nature doesn't provide exact measures. It provides perhaps exact counts, like "there are three peaches in this basket", but it doesn't provide exact measures like "this railing is 1 meter high". Furthermore our measures can vary a bit without affecting our reasonings or predictions, so it took a long time for scientists to see the problem. Usually there is no problem – with one exception: chaotic systems.

Some scientists had begun to be aware of the chaotic character of certain systems at the end of the XIXth century – for instance Henri Poincaré (1854 – 1912) when he studied the stability of the Solar system – and the problems it raised. In 1963, Edward Lorenz (1917 - 2008), an American meteorologist at MIT, published a startling discovery. He was studying numerically, with an early computer, the behavior of a set of differential equations linking a few variables describing the motion of the atmosphere. Every morning he would enter into the computer the numbers spit out at the end the previous day, in order to let the computer continue to crunch all day, thus pursuing the calculations of a trajectory. He noticed that if he changed a tiny bit (just rounding the numbers more or less) the input data, that would change dramatically the trajectory. He had found a chaotic system in a remarkably simple set of equations.

The harmonic oscillator, in phase space, just moves around in a circle with constant frequency, see figure 15, see also volume 1 of the collection *The Theoretical Minimum*, on classical mechanics, and more specifically chapter 8 on hamiltonian mechanics and time-translation invariance.

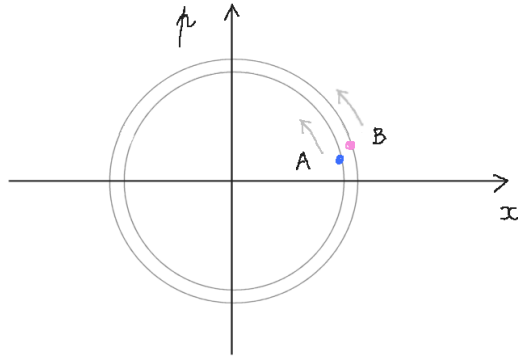


Figure 15: Harmonic oscillator. Representation in phase space. This system is a very non-chaotic. If we start from close by initial conditions, the trajectories are close to each other and remain so forever.

If we start from initial conditions A in phase space, the trajectory will be the inner circle. If we start from B , the trajectory will be the outer circle. The two trajectories are close to each other.

Going to a slightly more complicated system than the simple harmonic oscillator, a particle in an orbit around the Sun, in a true inverse square law, where we do exact Newtonian physics, is not chaotic. Two close by orbits will remain close by basically forever.

Exercise 1 : In the simple harmonic oscillator, whose phase space is represented in figure 15, do the two points A and B remain close to each other as time passes?

What about the case of a particle orbiting the Sun?

The three-body problem is chaotic in general. Neighboring orbits will eventually depart from each other.

In that case, predictability is a function of how accurately we know the initial conditions. To predict satisfactorily for a length of time t , we must know the initial conditions to within $\epsilon(t)$. And we have

$$\lim_{t \rightarrow +\infty} \epsilon(t) = 0 \quad (55)$$

So, in real life situations, necessarily at some point in terms of t , $\epsilon(t)$ will be smaller than the precision our instruments can achieve.

That is what it means to be chaotic.

Q.: Don't you also need to know the laws of motion with high precision?

A.: Yes, you also need to know the laws of motion with high precision, both the laws and the initial conditions.

Q.: What is the difference between the three-body problem and the two-body problem?

A.: One is chaotic and one is not. It is a mathematical property.

In our course in classical mechanics, we have seen the simple pendulum and the double pendulum.

The simple pendulum is not chaotic. It is very predictable.

But the double pendulum is chaotic. Suppose we start from two slightly different initial conditions, figure 16.

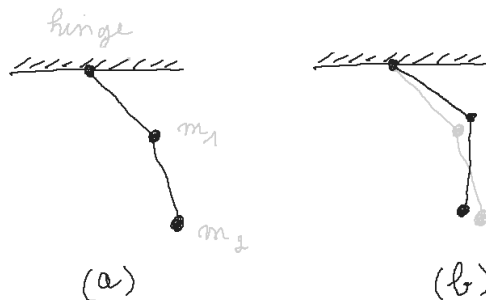


Figure 16: Double pendulum in two slightly different initial conditions.

We can release the double pendulum in configuration (a), or in configuration (b) – not even bothering to give it any initial momentum –, then it will almost immediately behave *quite differently*.

Have you ever seen a double pendulum in action? It is quite fascinating. You can follow it. After a little bit of time, it

will do something really weird. For instance, at some time t , all of the kinetic energy may unexpectedly go into one of the masses. Then the two will move again, not necessarily in the directions you expected. Very chaotic things happen.

Q.: How do we see it mathematically?

A.: Well, in volume 1, we wrote the equations of motion of the double pendulum. It was an example where the Lagrangian approach was much simpler than the Newtonian approach to obtain the equations.

Once we solve them, with numerical methods, we can represent in the phase space the trajectory depending on the initial conditions. We see that two trajectories starting from close by phase points depart radically.

But mathematically speaking, you are asking a very hard question. Why are some systems chaotic and other systems aren't chaotic? And can we readily distinguish them from some general characteristics of their equations?

I would say the state-of-the-art is still primitive. It is very hard to look at a hamiltonian and know whether it is chaotic or not, except for the fact that if it is complicated it is usually chaotic.

There are complicated systems, however, which are not chaotic. And, in 1963, the meteorologist Edward Lorenz found one which was not particularly complicated and was chaotic.

So this is not an easy question. It is something that mathematicians and mathematical physicists study. They don't have answers that are very general.

Q.: For the double pendulum to be chaotic, we need no friction, don't we?

A.: Right, the double pendulum with no friction.

Q.: So if we knew the initial conditions with infinite precision, there would be no chaos?

A.: If we knew the initial conditions to infinite precision – whatever that means – we could predict the motion to infinite time. But I think that is probably a meaningless statement.

The right statement is: There is a connection between the time period over which you want to predict the behavior of your system and how precise you have to know the initial conditions.

That has to do with how fast the trajectories in figure 13 depart from each other. If the trajectories tend to stay close to each other for a long period of time, the system is rather predictable for a long period of time. But eventually they will separate. So most systems are unpredictable in that sense.

Another way to say it is: Given any length of time that you want to predict, there exists a tolerance or a precision ϵ , such that, if you knew the phase point within that precision ϵ , it would allow you to predict for that length of time. You know, it is the kind of reasoning: Whatever t , there exists an ϵ such that and so forth and so on – the kind of mind twister some maths teachers in high school pretend to teach you maths with.

There is a theory by Lyapunov²¹, developed well before chaos had become an active domain of research in mathematics, which studies the stability of motions, and which has also relevant things to say about chaotic equations. The Lyapunov coefficient is the coefficient of the exponential growth of the separation between trajectories.

Q.: Is saying that a system is chaotic the same as saying that its equations don't have a closed solution?

A.: Chaotic systems certainly don't have closed solutions to the equations. But the converse is not true. It is possible to have a system which has no closed solution, and yet which is not chaotic.

Q.: What is the relationship between being unstable and being chaotic?

²¹Aleksandr Mikhailovich Lyapunov (1857 - 1918), Russian mathematician.

A.: Chaotic systems may or may not have unstable points. And an unstable system is not necessarily chaotic.

As we saw, the double pendulum is chaotic. If we think of one which can rotate 360° around its fixed hinge, it also has an unstable point. If you release it at rest with the two sticks vertical above the hinge, it is an unstable position. In theory it will stay there.

But if you give it the slightest perturbation, it will leave that state and not come back. Furthermore the stick attached to the fixed hinge can start to rotate on the right or on the left. You can also think of a pendulum with complete spherical possible rotation if you want to enrich the example.

But being chaotic is a different than being unstable. It means that from any small blob in the phase space trajectories depart quickly.

There are plenty of systems which have unstable states, but which are not chaotic. Orbits in the two-body problem are unstable, but they are not chaotic.

The simplest kind of instability is illustrated in figure 17. There is just one point of instability. It is at the very top of the hill. And, if you move the marble ever so slightly from that point, it can roll down in any direction. So from a small blob at that point in the phase space (position at the top, and zero velocity), the possible trajectories do depart from each other.

But from a blob around any other point in the phase space of that system, possible trajectories don't depart from each other, and are predictable over a long time unrelated to the precision of the initial conditions.

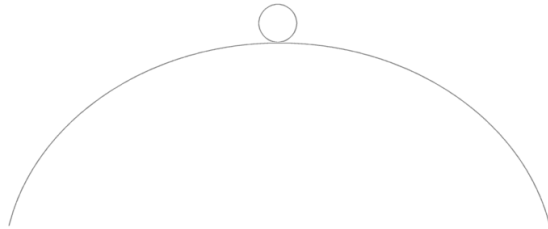


Figure 17: Unstable point in a non-chaotic system. A small perturbation on the marble will let it roll down the hill. And the possible trajectories from there are quite different.

The double pendulum, with a 360° fixed hinge, is of a different nature. It has one unstable rest point, the two sticks vertical above the fixed hinge, and one stable rest point, the two sticks hanging vertical below the fixed hinge. And it is chaotic everywhere, meaning that from any small blob around any point, the trajectories will quickly be different.

You can call any point – except the stable rest point – unstable if you like, although some people prefer to talk of an unstable point if, when perturbed the system won't come back to it, but which it won't leave spontaneously.

On the other hand a point is called stable if, when moving a bit away from it, the system comes back to it, whatever trajectory it follows to return.

So being unstable is different than being chaotic, which has to do with the behavior of trajectories leaving small blobs anywhere in the phase space. In a chaotic system, trajectories starting from any pair of nearby points in phase space depart radically from each other after some time.