


IP326. Lecture 1. Thursday, Jan. 3, 2019

- What is time-dependent statistical mechanics?

As the name suggests, time-dependent statistical mechanics is a generalization of time-*independent* statistical mechanics in which time is an additional variable. But this generalization is not a simple matter of introducing time dependence into the various functions that appear in time-independent statistical mechanics, and we'll spend rest of the lecture trying to understand what it does entail. And to do that, we'll first need to recall some facts about time-independent statistical mechanics.

Broadly speaking, time-independent statistical mechanics is the study of the connection between the macroscopic properties of matter *in equilibrium* and the microscopic properties of its constituent elements. The operative words here are: *macroscopic*, *microscopic*, and *equilibrium*, and we'll need to be clear about what these words connote before moving on. The following example will be helpful:

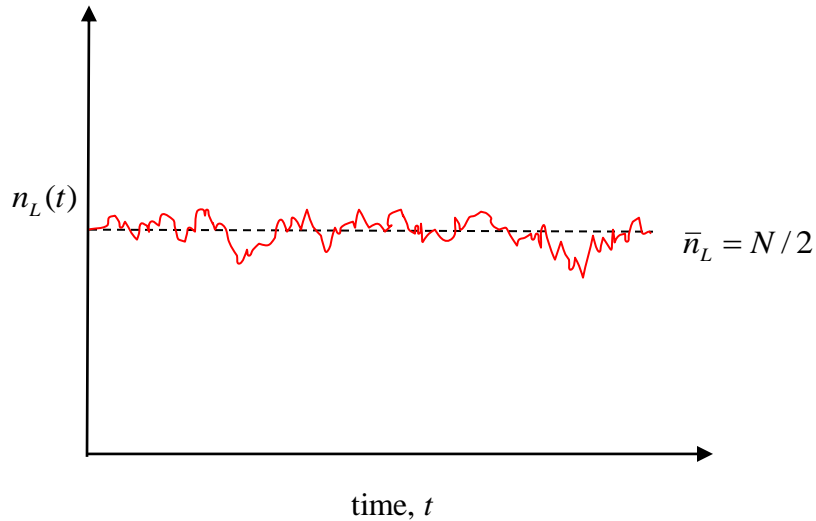
- Consider a gas of N independent *distinguishable* non-interacting particles enclosed in a sealed container, and imagine mentally dividing the container into two equal halves. A snapshot of the particles in this arrangement might look like this:



where the dots represent the positions of individual gas particles. (You have to imagine that these dots are actually distributed *randomly* throughout the box.)

Looking at these positions, we can obviously say that every particle is located in either the left half of the box or the right half. In a sense, then, every particle can be thought of as existing in two states – left or right. And one way of saying something about the state of the gas at a given instant of time is to say something about *every* particle in the box. Another way is to talk about the *number* of such particles in the left half of the box, without worrying about which one is located there. The first description applies at a *microscopic* level and the second at a *macroscopic* level, and in each the state

of the system varies with time because the particles are constantly moving from one side to the other through collisions with themselves or with the walls of the box. However, on *average*, if N is large enough, the *number* of gas molecules on the left is just $N/2$, and it doesn't change at all. In other words, if we were to track this number over time, this is what we would see, in graphical form



where $n_L(t)$ is the number of particles in the left half of the box at time t , with \bar{n}_L , the *average* of this number, which will be approximately $N/2$. How do we arrive at this average? Experimentally, in the following way: by taking snapshots of the gas at a series of distinct times; making a note of the number of molecules in the left half of the box at these times; adding the numbers together and dividing by the number of snapshots. That is,

$$\bar{n}_L = \frac{n_L(t_1) + n_L(t_2) + \dots + n_L(t_n)}{n} \approx \frac{N}{2}$$

and it corresponds to what we call a *time* average. If the time interval over which these measurements of particle number are made is large enough, and the intervals between successive measurements small enough, this average effectively becomes

$$\bar{n}_L(t_0, \tau) = \frac{1}{\tau} \int_{t_0}^{t_0+\tau} dt n_L(t)$$

where t_0 is the time at which the measurements are started and τ is the interval over they're made. The RHS of this equation is really only meaningful when τ is much greater than the time it takes for the system to rearrange itself. Ideally, we'd like τ to be infinitely large, in which case


$$\bar{n}_L(t_0) = \lim_{\tau \rightarrow \infty} \frac{1}{\tau} \int_{t_0}^{t_0+\tau} dt n_L(t)$$

independent of τ . And if additionally, the system is in equilibrium, it won't matter when the measurement is started, and there's then no dependence on t_0 either, which can be set to 0. Hence,

$$\bar{n}_L = \lim_{\tau \rightarrow \infty} \frac{1}{\tau} \int_0^\tau dt n_L(t)$$

- Time averages are what we tend to obtain automatically when we make experimental measurements of some macroscopic quantity because such measurements usually always take much, much longer than the time it takes for individual atoms or molecules to adopt new configurations. During the course of an experimental measurement, in other words, the system will have had more than enough time to make numerous transitions between different microscopic arrangements.

We could have thought about this average somewhat differently. Instead of recording the values of $n_L(t)$ at separate instants of time, we could have started with a large collection, say, M , of distinct containers of the same gas, all of them prepared with the same number of molecules under the same *macroscopic* conditions, and each divided into two halves by an imaginary partition. At some common instant of time, the situation, pictorially, is like the one below for a gas of 10 molecules.



Some boxes have less than $N/2$ molecules in the left half, some have more, but if you added them together and divided by M , you'd expect the result to be approximately $N/2$,

so long as both N and M were large. This way of calculating the average number of molecules in the left half of the box is called an *ensemble* average, and we'll denote it $\langle n_L \rangle$. So, in other words

$$\langle n_L \rangle = \frac{n_L^{(1)} + n_L^{(2)} + n_L^{(3)} + \cdots + n_L^{(N-1)} + n_L^{(M)}}{M}$$

The factor of $1/M$ in the above expression can be thought of as a weight factor, the fraction of replicas in the ensemble that have a given microscopic arrangement. In this example, each replica has the same weight, meaning, in effect, that there is no preference for any particular arrangement. Being a weighting factor, $1/M$ can also be interpreted as a probability of occurrence of that arrangement. So the ensemble average can also be more generally written as

$$\langle n_L \rangle = \sum_{i=1}^M p^{(i)} n_L^{(i)}, \quad p^{(i)} = \frac{1}{M}, \quad \forall i \in M \text{ at the same time } t$$

Ensemble averages are what we typically perform when we determine averages via statistical mechanics, and intuition suggests that there should be no difference between the time average of n_L and its ensemble average, and now we'll actually assert, as a **fundamental postulate** that

$$\bar{n}_L = \langle n_L \rangle \quad (\text{ergodic hypothesis})$$

- In the above illustrative example, we can think of n_L as a descriptor of the *macroscopic* state of the gas, that is, a state of the gas that makes no reference at all to anything about the gas molecules that it's made up of. Such a state is called a *macrostate* of the gas, and we regard the gas to be in **equilibrium** if the property that characterizes its macrostate completely doesn't change with time. Here, the variable n_L is the relevant property, and to all intents and purposes it is time-independent, because $n_L \approx \bar{n}_L \approx N/2$; in other words, if n_L were measured today, and then again tomorrow, and then yet again, say, a year from now, there would be no change at all in its value.

But as this example shows, equilibrium as we've defined it is not a static condition; it's actually a dynamic one; things are constantly happening at the *microscopic* level: particles are moving from one side of the box to the other; they're bouncing off the sides of the container; they're colliding with each other and exchanging energy. So at this level, the state of the gas – what we call the *microstate* – is constantly changing. And just as we characterized the macrostate of the gas by the variable n_L , we can identify an appropriate set of variables that characterize its microstate. Recall that in this example, every particle had just two states – it could either be in the left half of the box or the right. So one way to specify the gas's microstate is to specify the state – left or right – of every

molecule in it. So at these two levels of description, macro and micro, the associated states are the following:

Macrostate: n_L

Microstate: $\{1_\alpha, 2_\alpha, 3_\alpha, \dots, n_\alpha\}$, $\alpha = L \text{ or } R$, $t = t_1$

At equilibrium, $n_L \approx \bar{n}_L$ is a constant, even though the microstate changes from instant to instant.

So, to repeat, a macrostate of a system is some set of properties that characterizes the system in the bulk, without reference to the atoms or molecules it's made up of, a microstate of the system is a specification of some property of every single atom or molecule in the system at some instant of time, and equilibrium is the condition of time independence of the macrostate.

• Digression: some terminology. We'll need to decide what properties must be specified to identify macro- and microstates unambiguously. To do that, it will help to make a quick detour into some terminology, so that we know precisely what we mean by certain terms that we'll use over and over again in statistical mechanical contexts. These terms may seem familiar, but they are to be understood in a definite way. Let's look at some of them: (see Castellan, *Physical Chemistry*.)

System. Any part of the physical universe we're interested in studying.

Surroundings. That part of the universe lying outside the system.

Boundary. The physical surface enclosing the system and separating it from the surroundings. Boundaries can be of several kinds:

- Rigid or movable,
- Permeable (to matter flow) or impermeable or semi-permeable,
- Adiabatic (preventing heat flow) or diathermal (permitting heat flow).

Internal constraint. Any boundary within a system that divides it into two or more parts, called *sub-systems*. A system divided in this way into sub-systems is said to be *composite*. A system without internal constraints is said to be *simple*.

Closed system. One that is separated from the surroundings by a boundary that prevents any interaction between the two (Callen's definition; see his *Thermodynamics and an Introduction to Thermostatistics*). A closed system is completely isolated from its surroundings, and produces no observable effect or disturbance on it. (N.B. Other authors define a closed system as one in which no matter can pass through the boundary, though energy can, and an open system as one in which both energy and matter can pass through boundary.)

A system is said to have *properties*, which are those physical attributes that are perceived by the senses or can be made perceptible by certain methods of measurement. Different kinds of properties can be distinguished:

- Non-measurable, as the kinds of substances composing a system and the states of aggregation of its parts,
- measurable, to which a number can be assigned,
- *extensive*, which change in direct proportion to the size of the system,
- *intensive*, which are independent of the size of the system


A *phase* is a region within a simple system throughout which all of its properties are uniform.




A system is said to be in a definite *state* when each of its properties has a definite value.

A change in state of a system occurs when at least one of its properties changes its value. The change in state is completely specified when the initial and final states are specified. Systems initially in equilibrium will not change their state unless internal constraints are removed, or they are otherwise made to interact with their surroundings.

During a change of state, a system traverses a definite *path*, which is the ordered sequence of intermediate states lying between the initial and final states. If all the intermediate states in the path are equilibrium states, the path is said to be *quasi-static*. Physically, quasi-static means that if at any instant of time during the course of the change in state the system were isolated from its surroundings, all properties would be fixed at the values they had just prior to isolation. The path is said to be *reversible* if after the change of state, both system AND surroundings are returned to their initial states. Reversible paths are idealizations, but can be approximated by quasi-static paths.

● With these preliminaries, we can start looking at more realistic situations than the one we considered for illustrative purposes. And the first thing we'd like to do is identify the variables that we think are sufficient to completely specify the state of a system at both the macroscopic and microscopic levels. For the moment, we'll restrict our attention to *classical* systems that are simple, that are not translating or rotating, and that are not subject to external fields. Such systems can be placed under a range of different conditions that are relevant to experiment. Here are some common possibilities, shown schematically





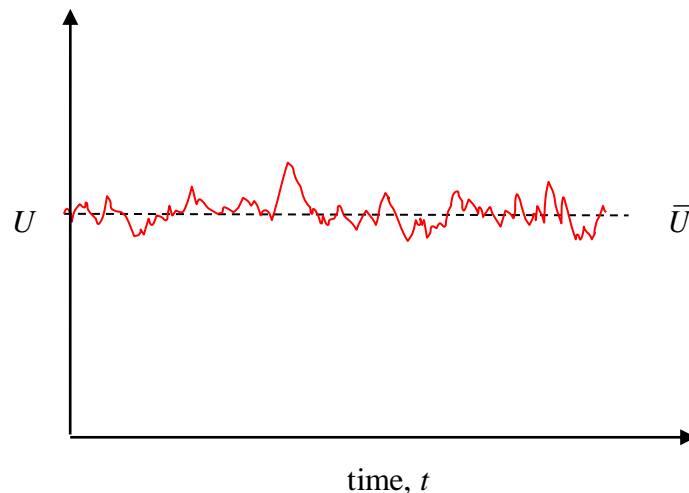
The first system is a *closed* system, completely isolated from its surroundings by a rigid, adiabatic, immovable boundary (the double lines). We'll state as a postulate (see Callen) that its macrostate is completely specified by U, V, N . (U is its internal energy – all the energy it possesses by virtue of the motion of molecules and their interactions with each other.)

The second is a *closed composite* system. The system of interest is the inner container; it's assumed to interact with a heat bath or thermal reservoir. A heat bath is any system so large or so designed that no matter how much heat flows into or out it, its temperature T remains the same. Here, the system and reservoir are separated by a diathermal boundary, so heat flows freely between the two. The macrostate of the system is T, V and N .

The third is a *closed composite* system. The system of interest is the inner container; its surroundings constitute a temperature and pressure reservoir. A pressure reservoir is a system so large or so designed that no matter how much its volume changes, its pressure stays the same. The system and surroundings are separated by a diathermal, movable boundary. The macrostate of the system is P, T, N .

The last is a *closed composite* system. The system of interest is the inner container; its surroundings constitute a temperature and particle reservoir. A particle reservoir is a system so large or so designed that no matter how many particles enter or leave it, its chemical potential remains the same. The system and surrounding are separated by a diathermal semi-permeable boundary. The macrostate state of the system is T, V, μ .

Apart from the first kind of system, which by itself is not very interesting since nothing happens to it, the others are slightly more complex realizations of our first illustrative example, in the following sense. Consider the system at T, V, N . Because the boundary is diathermal and energy U is freely exchanged between system and reservoir, the system has different amounts of energy at different instants of time, and a graph of U vs t would look like this:



As in the case of the gas, the variable U constantly changes from instant to instant, but the changes typically are very small if the system is large, and its average value, \bar{U} , obtained by adding up the energy values at different times and dividing by the number of measurements, is time-independent. Again, by virtue of this time independence of U (and the other macro variables), we say that the system is in equilibrium. An experimental measurement of U , because it takes longer than atomic/molecular timescales, is a time average of this variable.

A similar situation obtains for the system at constant T, V, μ . Now both energy and matter can flow across the boundary, which is diathermal and semi-permeable. So in this system, the macroscopic variables U and N fluctuate in time about an average value, which is the value that is recorded when that variable is measured physically, in a laboratory experiment. This average is time independent, and that being the case, the system is in equilibrium.

All of what we've just described refers to the macrostates of these systems. What about their microstates? We shall **postulate** that at the microscopic level, the system is completely specified by specifying the positions and momenta of every single particle in the system. If these variables are denoted q and p , we can now make the following identifications:


Macrostate: (U, V, N) or (T, V, N) or (T, P, N) .

Microstate: $\{q_1, p_1, q_2, p_2, \dots, q_N, p_N\}$

Macrostate: (T, V, μ)

Microstate: $\{q_1, p_1, q_2, p_2, \dots, q_i, p_i\}, \quad i = 1, 2, 3, \dots$

- Phase space. It's customary to think of these positions and momenta as spanning what we call a phase space, which for a collection of N particles in 3 dimensions is an abstract space made up of $6N$ mutually perpendicular axes, each axis corresponding to a particular Cartesian component of q or p for a particular particle. A point in this space is one possible microstate of the system. So, for example, if our system consisted of just a single particle in 1 dimension, its phase space would look like this:



p and q are called phase space variables, and as a matter of notation, we'll use the symbol Γ to represent the collection of phase space variables that define a particular microstate. That is, $\{q_1, p_1, q_2, p_2, \dots, q_N, p_N\} \equiv \Gamma$. (In $3d$, the p 's and q 's become vectors.)

Quantities that we measure experimentally are in general functions of phase space variables, and we refer to them as *dynamical variables*. They are fundamentally mechanical (as opposed to statistical, like a distribution), and they depend, in general on time, because the p 's and q 's depend on time. So if A denotes the dynamical variable in question, we have

$$A(\Gamma(t)) = A(q_1(t), p_1(t), \dots, q_N(t), p_N(t))$$

which for notational simplicity we often simply write as $A(t)$.

Examples of dynamical variables:

- Velocity of i th particle, \mathbf{v}_i
- Kinetic energy, or potential energy, or total energy
- Distance between particles i and j , $|\mathbf{r}_i - \mathbf{r}_j|$
- Density of particles at some point, $\rho(\mathbf{r}) = \sum_{i=1}^N \delta(\mathbf{r} - \mathbf{r}_i)$