1 Classical statistical physics

In the lectures we discussed the foundations of statistical physics within the framework of quantum mechanics. These brief notes are designed to introduce the basic ideas within the framework of classical physics.

1.1 Phase space density

Consider a system consisting of many particles e.g. a gas in a box. We will refer to a *microstate* of the system as a precise microscopic specification of the state of the system (i.e., the exact positions and velocities of each molecure). Of course this is very complicated! Fortunately, we are not usually interested in the details of exactly where each molecule is - we are interested in studying the *macroscopic* behaviour of the system. A *macrostate* corresponds to a macroscopic description of the system in terms of a few experimentally accessible parameters e.g. pressure, energy, volume etc. The goal of statistical physics is to predict properties of macrostates starting from the underlying microscopic laws of physics.

We want to predict properties of the system that don't depend on the details of the microstate, so we need to introduce some way of "averaging over microstates". This is done using an *ensemble* of systems. This is a set of $W \gg 1$ (noninteracting) copies of the system. Each system in the ensemble is specified by a particular microstate, which evolves according to the microscopic laws of physics of the system. Different systems in the ensemble can be in different microstates. What we'd like to do is to construct an ensemble for which taking an average over the ensemble lets us predict macroscopic quantities in physically relevant situations.

To specify the microstate of a system we will work in *phase space*. If you've taken *Classical Dynamics* you will know what phase space is. Otherwise here is a quick summary. For a single particle, phase space in the 6-dimensional space with coordinates $(\mathbf{q}, \mathbf{p})$ where $\mathbf{q}$ is the particle's position and $\mathbf{p}$ is its momentum. As the particle moves, it traces out a unique trajectory in phase space. For a system containing N particles, phase space is the $6N$-dimensional space with coordinates $(\mathbf{q}_1, \ldots, \mathbf{q}_N, \mathbf{p}_1, \ldots, \mathbf{p}_N)$. The microstate of the system at any time is specified by a single point in phase space, and this point traces out a trajectory over time. Different initial conditions give different trajectories. These trajectories can't intersect because that would contradict the fact that there is a unique trajectory through each point.

Now consider an ensemble of systems. At a given time, each system in our ensemble is fully specified by a single point in phase space, so the ensemble corresponds to W points in phase space. Since $W \gg 1$ we can think of these points as distributed in phase space with some density $\rho(\mathbf{q}_i, \mathbf{p}_i)$. By this we mean that, the number of systems occupying a region of phase space centred on the point $(\mathbf{q}_1, \ldots, \mathbf{q}_N, \mathbf{p}_1, \ldots, \mathbf{p}_N)$ with volume $d^{3N}\mathbf{q}d^{3N}\mathbf{p} = d^3\mathbf{q}_1 \ldots d^3\mathbf{q}_N d^3\mathbf{p}_1 \ldots d^3\mathbf{p}_N$ is given by

$$W\rho(\mathbf{q}_1, \ldots, \mathbf{q}_N, \mathbf{p}_1, \ldots, \mathbf{p}_N) d^{3N}\mathbf{q}d^{3N}\mathbf{p} \tag{1.1}$$

It may help to visualize the points describing the ensemble as a fluid in phase space with density $W\rho$. From this definition, the phase space density ρ satisfies the normalization condition

$$\int \rho d^{3N}\mathbf{q}\ d^{3N}\mathbf{p} = 1 \tag{1.2}$$

We do not assume our macroscopic system to be in some particular microstate. Instead we can talk about the *probability* of it being in some microstate. More precisely, this is the probability that a system picked at random from the ensemble will be in this microstate. To calculate this probability we interpret ρ as a *probability density* in phase space:

The probability that the actual microstate of our system is lies inside a phase space volume $d^{3N}\mathbf{q}\ d^{3N}\mathbf{p}$ *centred on the point* $(\mathbf{q}_i, \mathbf{p}_i)$ *is given by* $\rho(\mathbf{q}_i, \mathbf{p}_i) d^{3N}\mathbf{q}\ d^{3N}\mathbf{p}$*.*

Now we turn to dynamics: we need to work out how $\rho(t, \mathbf{q}_i, \mathbf{p}_i)$ depends on t. We assume that our system satisfies an equation of motion arising from a Lagrangian. In the *Classical Dynamics* course it is shown how one can rewrite the equations of motion in terms of a Hamiltonian, and the resulting equations of motion translate into a simple equation of motion for ρ:

$$0 = \frac{d\rho}{dt} \equiv \frac{\partial \rho}{\partial t} + \sum_i \left(\frac{d\mathbf{q}_i}{dt} \cdot \nabla_{\mathbf{q}_i} \rho + \frac{d\mathbf{p}_i}{dt} \cdot \nabla_{\mathbf{p}_i} \rho \right) \tag{1.3}$$

where on the RHS we evaluate $d\mathbf{q}_i/dt$ and $d\mathbf{p}_i/dt$ using the equations of motion. Just as in fluid mechanics, $d\rho/dt$ is simply the derivative moving along a trajectory in phase space. Hence this equation has a very simple interpretation: *ρ is constant in time following the flow in phase space.* This is *Liouville's theorem.*

Consider a region of phase space small enough that ρ is constant throughout the region. Let the region have volume $dV = d^{3N}\mathbf{q}\ d^{3N}\mathbf{p}$, so it is occupied by $n = W\rho dV$ systems of our ensemble. Under time evolution, these systems will evolve to occupy some other region of phase space, with volume dV'. Since ρ is constant along the flow we have $n = W\rho dV'$, so it follows that $dV' = dV$. In other words, *time evolution preserves volume in phase space.* However, the *shape* of this region will change, and can become *very* complicated.

1.2 Statistical equilibrium, the fundamental assumption and the microcanonical ensemble

To calculate anything we need to know what ρ is. In this course we will be entirely concerned with systems which are in *statistical equilibrium.* By this we mean that probabilities are not changing with time. Since ρ is a probability density function, this implies that ρ cannot be varying with time, i.e., $\partial\rho/\partial t = 0$ or $\rho = f(\mathbf{q}_i, \mathbf{p}_i)$ for some function f. The equation of motion $d\rho/dt = 0$ now tells us that $f(\mathbf{q}_i(t), \mathbf{p}_i(t))$ is independent of t, i.e., it is *conserved* along any trajectory. Now most dynamical systems admit just a few conservation laws such as conservation of the total energy, total momentum and total angular momentum. So in statistical equilibrium ρ must be some function of these conserved quantities.

We will often be interested in systems such as a gas in a box where total momentum and total angular momentum are not conserved because of interactions with the walls of the box. So in practice the only conserved quantity at our disposal is the total energy E and so $\rho = f(E)$ in statistical equilibrium.

We have seen that, in statistical equilibrium, the problem of specifying a function ρ on a $6N$-dimensional phase space reduces to specifying a single function $f(E)$. But what dictates our choice of f? By observing our macroscopic system we will have some partial knowledge about the possible microstates that it could be in. In particular, we will have some idea about its energy. Let's assume that we know that the energy of our system is known to lie between E and $E + \delta E$ where δE represents our experimental uncertainty. We would then choose f to vanish outside this range. But how do we choose f inside this range?

The *fundamental assumption* of classical statistical physics is

All microstates consistent with our partial knowledge of the system are equally probably in phase space.

In other words, we must take ρ to be constant in the region of phase space corresponding to energy between E and $E + \delta E$, and zero outside this region. Stating that the energy is E defines some surface in phase space, so the region of interest lies between the surfaces corresponding to energy E and $E + \delta E$. Let $\Omega(E, \delta E)$ be the phase space volume of this region. Then from the normalization condition (1.2) we have

$$\rho = \begin{cases} \frac{1}{\Omega(E,\delta E)} & \text{if energy between } E \text{ and } E + \delta E \\ 0 & \text{otherwise} \end{cases} \tag{1.4}$$

This defines the *microcanonical ensemble.*

As the name suggests, the fundamental assumption cannot be derived from something more basic. It is an axiom of statistical physics. Note that it refers to equal probabilities in *phase space*, not in energy space e.g. it is $\rho(\mathbf{q}_i, \mathbf{p}_i)$ that is constant, not $f(E)$. Why should we make this assumption? Because we need to assume something about ρ to make predictions, this is the simplest assumption, and it works!