Statistical Physics

Matthew Wingate[*]

14 January 2010

---

[*] Comment & corrections to `M.Wingate@damtp.cam.ac.uk` . Notes and other information also available at `http://www.damtp.cam.ac.uk/user/wingate/StatPhys`

1 Introduction

1.1 Overview

In previous theoretical physics courses you learned how to describe the motion of a classical body like a billiard ball and how to find the quantum states of a microscopic body like a single atom or electron, perhaps under the influence of an external potential. You possibly have studied, or will study 2 bodies, classical or quantum, scattering off of one another. In most cases the interest was clearly exactly in those microscopic degrees-of-freedom.

The fields of thermodynamics and fluid dynamics, on the other hand, deal with macroscopic phenomena. They focuses on relationships between macroscopic observables like pressure, volume, and temperature. Many of these relationships are the same, or similar, for all systems, up to coefficients like heat capacities. These coefficients are properties of specific gases and cannot be computed with the laws of thermodynamics. They must be measured empirically, it would appear. In fact, thermodynamics was thought to be fundamental until the reality of atoms was accepted, an acceptance which came in large part due to the success of statistical mechanics.

Statistical physics is a bridge between the microscopic and the macroscopic. No matter what the individual degrees-of-freedom, classical or quantum, it provides a framework for studying a large number of assembled bodies. It takes as its input the Hamiltonian H of a system, defines something called a partition function Z, and allows us to compute thermodynamic relationships. Furthermore, specific coefficients are computable with statistical mechanics from microscopic first principles. In current research, the well-known laws governing nuclear interactions and quark interactions are being used to find ways to distinguish between neutron stars and quark stars. It is presently not known if the latter can exist. By making firm theoretical predictions perhaps a good signature of quark stars can be uncovered and then a search for this signature can be done.

The bridge can also be traversed in the other direction. If the microscopic behaviour, *i.e.* the Hamiltonian, is unknown then measurements of thermodynamic properties can constrain the microscopic properties. A classic example is the spectrum of black-body radiation. Classical derivations did not agree with the observed spectrum at low energies. It was only when Planck posited that light could be interpreted as discrete quanta that correct spectrum emerged from his calculations.

Figure 1 gives a picture of this path from microscopic to macroscopic. Statistical physics methods are also being extended to model financial markets and traffic patterns.

1.2 Microstates

A microstate is a unique arrangement of the microscopic degrees-of-freedom. Before delving into the basics of statistical mechanics, it is useful to have a few simple systems in mind. We will use these to illustrate key ideas throughout the course.

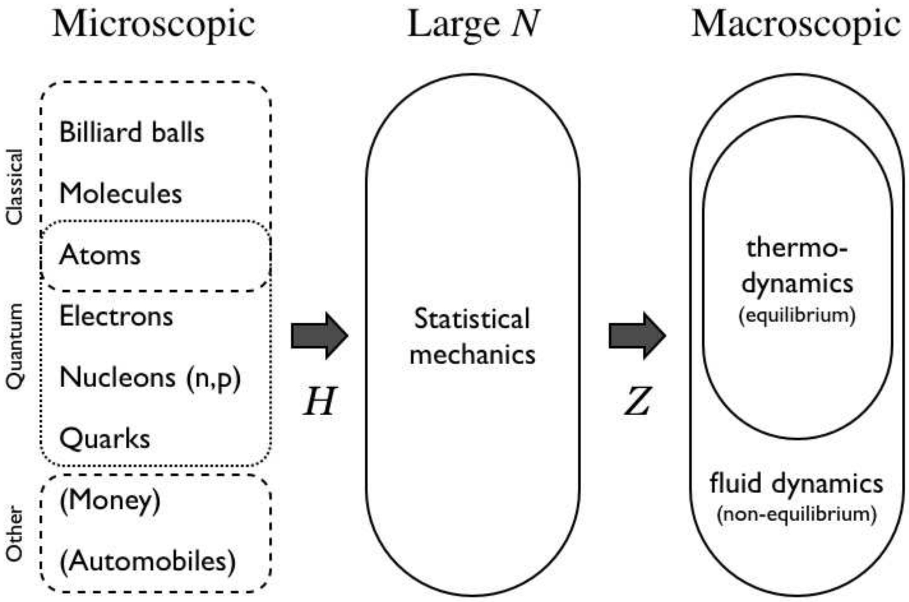

Figure 1: Statistical mechanics is a bridge between microscopic degrees-of-freedom and macroscopic behaviour. Given a Hamiltonian H, statistical physics defines a partition function Z from which thermodynamic (and hydrodynamic) properties may be calculated.

Example: spin model

Let us take a system of N particles with spin $\frac{1}{2}$, such that the particles are fixed in space. The only degree-of-freedom a particle has, let's say, is to be aligned in the $\pm z$-direction. Let's also assume an external magnetic field $\mathbf{B}$ points in the $+z$-direction and the energy of a particle is $\varepsilon_+ = -\mu B$ if it is aligned with the magnetic field and $\varepsilon_- = +\mu B$ if it is anti-aligned (μ is the particle's magnetic moment). The net energy E of the system is simply the sum of the spins' energies. Table 1 lists all microstates for the $N = 3$ model along with the energies of each microstate.

Example: non-interacting quantum particles

Let us recall an example from Part IB Quantum Mechanics. This will be a very useful one through our course. Consider the stationary states of a single particle in a cubic box with periodic boundary conditions: $V = L^3$ and $U(\mathbf{x}) = 0$. The stationary states are solutions to the time independent Schrödinger equation

$$-\frac{\hbar}{2m}\nabla^2\psi = \varepsilon\psi \tag{1.2.1}$$

which we recall to be plane waves, labelled by a wave vector $\mathbf{k}$

$$\psi_{\mathbf{k}}(\mathbf{x}) = \frac{1}{\sqrt{V}}e^{i\mathbf{k}\cdot\mathbf{x}} \tag{1.2.2}$$

State index	state	E
1	$\lvert +++\rangle$	$-3\mu B$
2	$\lvert ++-\rangle$	$-\mu B$
3	$\lvert +-+\rangle$	$-\mu B$
4	$\lvert -++\rangle$	$-\mu B$
5	$\lvert --+\rangle$	μB
6	$\lvert -+-\rangle$	μB
7	$\lvert +--\rangle$	μB
8	$\lvert ---\rangle$	$3\mu B$

Table 1: Microstates for $N = 3$ spin model.

where $|\mathbf{k}| = \sqrt{2m\varepsilon}/\hbar$. Periodic boundary conditions, $\psi(x_i + L) = \psi(x_i)$ imply the allowed values of k_i are discrete

$$\mathbf{k} = \frac{2\pi}{L}\mathbf{n}, \quad \mathbf{n} = (n_x, n_y, n_z)\,. \tag{1.2.3}$$

Single particle microstates can be labelled by the wave vector $|\mathbf{k}\rangle$. Since this is a noninteracting theory the N particle state is just the outer product of N single particle microstates

$$|i\rangle \equiv |\mathbf{k}_1\mathbf{k}_2\cdots\mathbf{k}_N\rangle \equiv |\mathbf{k}_1\rangle \otimes |\mathbf{k}_2\rangle \otimes \cdots \otimes |\mathbf{k}_N\rangle\,. \tag{1.2.4}$$

Granted, most of us do not live in periodic boxes. Actually any boundary conditions will yield a discrete spectrum: sinusoidal wavefunctions must have a discrete number of nodes. We will see that it is easier to work with a discrete spectrum of states and finite N rather than an infinity of particles each of which has a continuous energy spectrum. Usually at some point in the calculation we take the infinite particle, infinite volume limit, holding the particle density fixed:

$$N \to \infty, \quad V \to \infty, \quad \frac{N}{V} \text{ fixed} \tag{1.2.5}$$

This limit is called the **thermodynamic limit**.

Example: classical gas

We will not discuss the classical Hamiltonian again until Chap. 10. Nevertheless it is worth pointing out that classical statistical mechanics is logically well-formulated, despite the empirical observation that experiments disagree with classical predictions in the quantum regime, and a couple "paradoxes" can only be understood using quantum mechanics.

The classical Hamiltonian of a gas of N particles is a function of coordinates q_s and their conjugate momenta p_s, both of which can be known in principle to arbitrary accuracy. (s is an index which runs over the degrees-of-freedom. Noninteracting structureless particles in 3-dimensions are described by $3N$ coordinates and $3N$ momenta.) There are a continuous infinity of microstates represented by the phase space spanned by the coordinates and conjugate momenta.

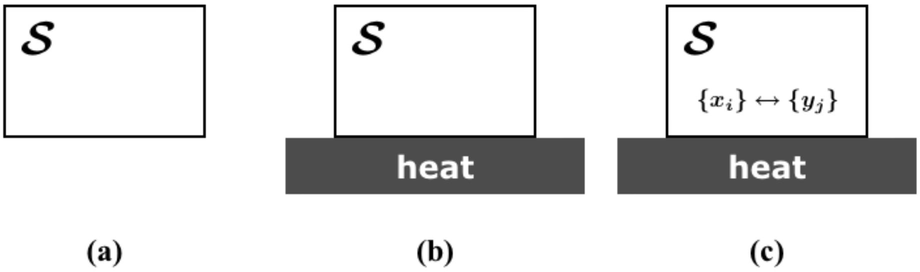

Figure 2: Schematic depiction of (a) the microcanonical ensemble, (b) the canonical ensemble, and (c) the grand canonical ensemble.

1.3 Ensembles and subsystems

Not all microstates will be accessible, or allowable in a given scenario. Suppose we take the spin model described above and assert that we are studying the $N = 3$ spin model which has $E = -\mu B$. Then the only 3 allowable states labelled 2-4: $|++-\rangle$, $|+-+\rangle$, and $|-++\rangle$.

The set of allowable microstates under a particular circumstance is termed an **ensemble**. Under different circumstances different ensembles are appropriate. There are three ensembles we will consider:

- **Microcanonical ensemble**

  In ensemble $\mathcal{E}_{\text{ME}}$, the total energy E and number of particles N is fixed. This is appropriate for a truly isolated system.

- **Canonical ensemble**

  $\mathcal{E}_{\text{CE}}$ has N particles fixed, but E can fluctuate about a mean $\langle E \rangle$ due to thermal contact with another system, sometimes called a heat bath.

- **Grand canonical ensemble**

  In $\mathcal{E}_{\text{GCE}}$ both E and N fluctuate about mean values $\langle E \rangle$ and $\langle N \rangle$. This is appropriate when a heat bath is present, as well as something varying particle number. For example, particle number is not conserved in thermal systems where a chemical reaction is occurring in both directions, or in the early universe when electrons, positrons, and photons were all in equilibrium.

These three ensembles are depicted in Figure 2. Recalling the example of free quantum particles in a periodic box, we can see the three ensembles correspond to different allowed microstates (1.2.4). In the microcanonical ensemble, all the $|i\rangle$ have N particles and only $N-1$ values of $|\mathbf{k}|$ are unconstrained, the last constrained by energy conservation. In the canonical ensemble, E is not conserved so we have one fewer constraint compared to the microcanonical ensemble. In the grand canonical ensemble, N can fluctuate, and correspondingly so do the single particle states contributing to $|i\rangle$.

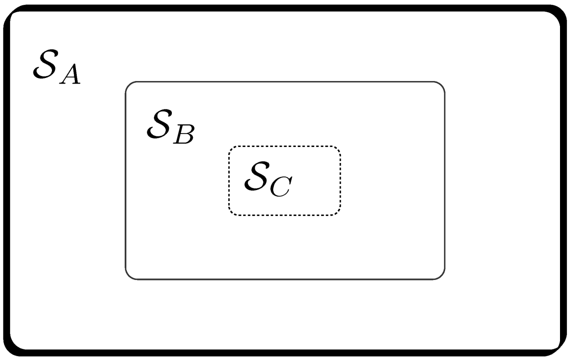

Figure 3: Describing different subsystems with different ensembles: System $\mathcal{S}_A$ is thermally isolated from the outside environment and described by the microcanonical ensemble. Subsystem $\mathcal{S}_B$ contains particles which can exchange energy with those outside, but are confined within $\mathcal{S}_B$, so the canonical ensemble is most relevant. Subsystem $\mathcal{S}_C$ is simply a volume within $\mathcal{S}_B$ through which energy and particles can enter or leave, so the grand canonical ensemble applies.

Another way to view these three different ensembles is depicted in Figure 3. A system $\mathcal{S}_A$ described by the microcanonical ensemble is completely isolated from any outside environment. A system $\mathcal{S}_B$ described by the canonical ensemble can be thought of as a subsystem, in thermal contact with a larger subsystem with which it can exchange small amounts of energy, but not permitting particles to move in or out of the subsystem. An example use of the grand canonical ensemble would be to describe a subvolume $\mathcal{S}_C$ of an open container: particles and energy are free to move in and out of the area.

The thermodynamic laws that emerge from the statistical mechanics of these different ensembles are identical. This is because the fluctuations about ensembles averages are so small. We will estimate the sizes of fluctuations throughout. However, the different ensembles give us flexibility in carrying out the statistical mechanics calculations, as we shall also see.

1.4 Equilibrium

Consider an isolated, or closed, system $\mathcal{S}$. No heat enters or exit, and there are no external forces. It is an empirical fact that after sufficient time, a steady state is reached which is independent of the initial condition. This steady state is characterised by a set of time-independent variables pressure, volume, temperature, entropy, energy, and so on. This state is called **equilibrium**. **Thermal equilibrium** specifically refers to a steady state with respect to heat transfer while **chemical equilibrium** refers to a steady state with respect to particle composition.

1.5 Ergodic hypothesis

We must make a crucial hypothesis, that of ergodicity: a system $\mathcal{S}$ evolves through *all* microstates belonging to the relevant ensemble eventually. That is, there is a mathematically nonzero probability that every microstate in the ensemble will be realised if we wait long enough. The set of probabilities ρ_i, where i labels the microstates, are independent of time for a system in equilibrium. In that case, it is no longer necessary to consider the order in which the system visits the assortment of microstates. We simply need to know these probabilities.

The consequence of making this hypothesis is a significant simplification. What we mean by expectation value is really a time average. For some observable X, its expectation value is defined to be

$$\langle X \rangle \;\equiv\; \lim_{T\to\infty} \frac{1}{T} \int_0^T X(t)\,\mathrm{d}t\,. \tag{1.5.1}$$

However, if we really needed to evaluate $X(t)$ for this whole time, we would be in trouble. We would need to solve the 10^{23}-body problem completely! The ergodic hypothesis means that we do not need to worry about how the system evolves in time. Since every microstate will eventually be visited, we can sum over all microstates, weighted by the relative frequency with which that microstate will be visited. That is, we can evaluate expectation values as ensemble averages, properly weighted averages over microstates:

$$\langle X \rangle \;=\; \sum_i \langle i|\hat{X}|i\rangle\, \rho(E_i) \tag{1.5.2}$$

where $\rho(E_i)$ is a **probability density** which dictates which microstates are accessible, and how accessible. We will discuss later how we know $\rho_i = \rho(E_i)$ is a function of energy only.

To repeat, since each microstate must appear at some time in the infinity of time, we can forget about the time history of the system and imagine an infinite set of **replicate systems**.

1.6 Fundamental postulate

In order to derive the probability densities for different ensembles we begin with an assumption, probably the most fundamental postulate of statistical mechanics

> *An isolated system in equilibrium is found in any of its accessible microstates with equal probability.*

Note that this assumption is for a completely isolated system, so is relevant for the microcanonical ensemble, but not the canonical or grand canonical ensembles.

For the microcanonical ensemble, where the system energy E is strictly conserved

$$\rho_i^{\mathrm{ME}} \propto \delta(E - E_i)\,. \tag{1.6.1}$$

All microstates which have $E_i = E$ are weighed equally, while those with $E \neq E_i$ are excluded. The microcanonical ensemble is quite useful, but in this short course we must move now to the more broadly useful canonical ensemble.

Further reading

D Lindley, *Boltzmann's Atom: the Great Debate that Launched a Revolution in Physics*, (The Free Press, 2001).