

# 1

## *Introduction: thermodynamic systems seen from outside*

### 1.1 Background ideas

This book is about thermodynamics, the study of heat. We shall start very near the beginning. I imagine that you, the reader, have some general knowledge of the subject. We all have an intuitive knowledge of hotness and coldness. Beyond that, I suppose that you are familiar with the idea of measuring hotness, or *temperature*, with a simple thermometer, some device whose properties are seen to change when it is placed in contact with a hotter or colder body. Given such a device – a mercury-in-glass thermometer, perhaps, in which the length of the mercury column increases with temperature – you know that we can set up an *empirical temperature scale* by assigning temperatures to two fixed points (as in the Celsius scale, where 0 °C is the temperature of melting ice and 100 °C is the temperature of boiling water), and then dividing the scale in some arbitrary way, by making temperature rise proportional to the increase in length of the mercury column, for instance.

Celsius temperature scale

Ideal gas temperature scale

You may have heard, too, of the *ideal gas scale of temperature*, defined by making the temperature proportional to the pressure exerted by a sample of gas of low density held in a container at fixed volume, and defining the triple-point temperature of water to be 273.16 K. (Ideal gas temperatures are given the symbol  $\theta$ , and measured in units of *degrees kelvin*, or K. The value for the temperature of the triple-point of water is chosen so that this unit is equal in magnitude to the unit of the Celsius scale.) You may know that the temperature 0 K is known as *absolute zero*, and that this is thought to be the coldest possible temperature, but that there is no limit to how hot a body may be.

Heat capacity

I assume also that you have met the idea that heat is a form of energy. We may, for instance, heat up a liquid by dissipating electrical energy in a heating coil. The heat required to do this may be measured electrically (in joules). The heat required to raise a body through unit temperature is called its *heat capacity*  $C$  and has units of  $\text{J K}^{-1}$ . Heat may also be needed to make a body *change phase* without changing the temperature – to melt ice, for instance. The heat required is known in this case as *latent heat*  $L$ .

Latent heat

2      *Introduction: thermodynamic systems seen from outside*      1.2

### 1.2      **The approach to thermodynamics through statistical mechanics**

In the eighteenth century some scientists thought that the phenomena of heat, like the phenomena of electricity, involved a new physical mechanism. Joseph Black, for instance, suggested that heat was a fluid known as ‘caloric’ which penetrated ordinary matter. At that time it seemed that the proper approach to the subject was to proceed empirically, to codify the experimental results into laws and to make deductions from those laws. We now know that this early approach was wrong in so far as it was looking for a new physical model. We have discovered, as Newton and Boyle had earlier suspected, that ‘hot’ bodies are simply bodies which contain a large amount of randomised energy, and that the atoms of which the bodies consist are perfectly ordinary atoms obeying the well-known laws of mechanics (or, rather, quantum mechanics). The laws of thermodynamics, therefore, though they first appeared as experimental laws, *ought to be deducible from the laws of mechanics*. Many textbooks still approach classical thermodynamics as an experimental, empirical subject, and the reasons are not wholly historical. The empirical approach to classical thermodynamics is logically very satisfying and, because it establishes rigorous links between thermodynamic facts without explaining the facts themselves in detail, it helps to emphasise the generality of the results in the student’s mind. But, in my experience, it also has bad effects. Students, though they see the subject as beautiful, may also find it mystifying. It feels strange to be handling concepts such as temperature and entropy, which are clearly important, without having any idea of what they mean at the atomic level. Moreover, those students who do go on to study the mechanistic interpretation – ‘statistical thermodynamics’ – usually find it hard to bridge the gap between the two approaches. In this book, therefore, the approach is mechanistic – we *start* with the mechanical picture and derive the laws of thermodynamics from the mechanics. This means that the early part of the presentation may seem a long way from the everyday world of experiment. But I do not think that the approach is harder to understand than the traditional one, and I believe that in the long run it pays handsome dividends. In this introductory chapter, before plunging into the details of the mechanistic picture, we make some general comments on the nature of thermodynamic systems, how they behave as seen from outside, and what the tasks of the mechanistic theory will be.

What, then, is a thermodynamic system, and in what way does thermodynamics differ from mechanics (or quantum mechanics)? The answer is simple. There is no *fundamental* difference between a mechanical system and a thermodynamic one, for the name just expresses a difference of attitude. By a thermodynamic system we mean a system in which *there are so many relevant degrees of freedom that we cannot possibly keep track of all of them*. We think of the solar system of the Sun and nine planets as a mechanical system because, using Newton’s laws of gravitation and mechanics we can in practice, knowing the starting positions and velocities of all the bodies in the system, make good

Thermodynamic systems

### 1.3 *The thermodynamic state*

3

predictions of the future motions of all the bodies. We think, on the other hand, of a box of helium gas containing  $10^{23}$  atoms as a thermodynamic system because, even if we had values for the starting positions and velocities for every atom, or the equivalent wave function for the gas, we could never in practice predict their future behaviours, partly because the calculations would be prohibitively long (even with the largest conceivable computer) and partly because we could never hope to have sufficiently accurate values of the starting conditions. The state of the helium gas is just as predictable *in principle* as the state of the solar system – we know the forces and the laws of motion. But in practice the prediction cannot be done because there are too many variables.

It follows that any theory of the gas has to be statistical in nature. We cannot speak of where a particular atom is, but we may speak of the *probability* of finding it in a given position and moving with a particular velocity. We shall examine the meaning of probability in thermodynamic systems more closely in the next chapter. Here we just note that our theory of the gas must proceed by the following route. Because we do not start with exact knowledge of the state of the gas, we have to assume, at some suitable starting point in time, a plausible or reasonable description of the system in terms of probabilities – known as the *initial probabilities*, the probabilities already present when we first meet the system. Then, from our knowledge of mechanics or quantum mechanics, which shows how the system would develop from each of its possible starting states, *we can find how the theoretical probabilities change with time*, and hence make predictions (which can themselves only be in terms of probabilities) of the properties of the system at future times. For instance, if the gas is known originally to have an uneven density, we represent this by a corresponding uneven probability distribution for the positions of the molecules, and, knowing how the molecules move, we can compute the probability distribution for their positions at later times. Knowing the theoretical distributions at later times, we can calculate how observable quantities, such as the pressure at a point in the gas, would be expected to behave on average. Because such predictions are essentially statistical, we can never make an *exact* prediction of the pressure. This is, however, a less serious limitation than might be expected. It turns out, as we shall see, that for many quantities the predicted statistical distributions are so sharply peaked that the statistical prediction is, in effect, an exact prediction for most practical purposes.

#### 1.3 The thermodynamic state

What do we mean when we say that some system is in a definite *thermodynamic state*? First of all, to specify the state at some particular instant we must fix the *constraints*. A constraint means in thermodynamics some parameter of the system which can be held fixed or varied at will by the observer. The volume of a gas and the electric charge on a capacitor are examples of constraints. We usually treat a constraint as having a definite and

§ 2.1

The statistical  
method in outline

Constraints

exact value which can be fixed from outside the system (though this is an idealisation).

For a given set of constraints, the system still has access to a very large number of *microscopic states* – for a gas held at fixed volume, for instance, the molecules could be started at time zero in many different positions and with many different velocities. As time proceeds and collisions occur in the gas the microscopic state is continually changing: we say that it is subject to *thermal fluctuations*.

Microstates

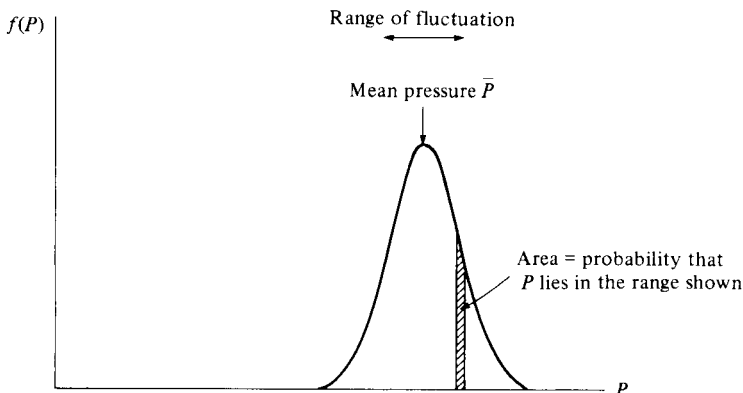
Fluctuations

For a given set of constraints there are also many properties of the system which can be measured – the *thermodynamic observables*. Some observables may be thought of as responses to particular constraints. The pressure, for instance, is the system's response to the volume constraint, and the voltage between the capacitor plates is the system's response to the charge on them. Other measurable quantities such as the refractive index or the X-ray diffraction pattern are less obviously related to particular constraints. These are all observables.

Observables

Since the measured value of an observable must depend on the microscopic state of the system, the values of the observables will also be subject to thermal fluctuations. If sampled many times over a short period all observables have a *fluctuation distribution* (Fig. 1.1). The observable will have a mean value, but will sometimes lie above, and sometimes below, the mean. Where the observable involves the combined effect of a large number of particles (as with the pressure of a gas on a large area) its fluctuation

Fig. 1.1. The distribution of fluctuations in the pressure  $P$  of a gas on a very small area. The pressure is due to random bombardment by gas molecules and therefore wanders around the mean value. Such distributions are described by a *distribution function*  $f(P)$ , which is defined so that the probability of finding  $P$  in the narrow range between  $P$  and  $P+dP$  is equal to the corresponding area under the curve,  $f(P) dP$ , which is shown shaded. The total area under the curve is unity. The mean pressure  $\bar{P}$  is  $\int P f(P) dP$ .



### 1.3 *The thermodynamic state*

5

distribution will normally be narrow and the fluctuations fairly insignificant. If only a few particles are involved (as with the density measured over a small volume) the fluctuations may be relatively large. We shall regard the system as being in a definite thermodynamic state when its observables take definite (but not necessarily equilibrium) values. In statistical terms this means that *in a definite thermodynamic state we know the probabilities of the various micro-states sufficiently accurately to make good predictions of the values of the observables*, averaged over their fluctuations.

Thermal equilibrium

We are often interested in systems in *thermal equilibrium*. What does this mean? In essence, we say that a system has reached thermal equilibrium when its observables have ceased to change with time. For instance, a gas whose molecules are all at one end of the container is not in thermal equilibrium because the density at that end of the container – an observable – will be falling with time. But such a description of equilibrium has to be interpreted cautiously. Because of the fluctuations, we cannot define equilibrium by saying that the observables are strictly constant. Instead, we have to say that *the system is in thermal equilibrium when the statistical distributions of the observables, averaged over their fluctuations, are constant*. In a few cases the time to reach equilibrium may be so long that we risk misinterpreting the situation. A piece of glass, for instance, may appear to be in thermal equilibrium. Actually, it is not; the true equilibrium state is crystalline, but glass has no crystal structure. It takes many thousands of years for glass to crystallise, and in this case even the fluctuations in the arrangement of atoms are so slow that we do not notice them. Strictly speaking, to define the *full* thermodynamic equilibrium state for glass we ought to average over the fluctuations for millions of years, watching all the time to see when the distribution, averaged over the very slow fluctuations, has itself settled down to a statistical equilibrium state.

What fixes the nature of the equilibrium state?

Once the system has reached equilibrium, the observables all take definite mean values. It is important to ask what fixes these equilibrium values, apart from the constraints. (We know, for instance, that if we put a given quantity of gas into a given volume it does not always show the same pressure in equilibrium.) The answer given by experiment is a remarkable one. We shall see in the next section that the *internal energy*, the total energy inside the system, is a well-defined quantity. What experiment almost always shows is that *the nature of the equilibrium state is fixed entirely by the constraints and the value of the internal energy*. Although the system may have started from a great variety of different and macroscopically distinct non-equilibrium states, most of the details of these states have no effect on the nature of the equilibrium state. Evidently the randomising process which goes on as equilibrium is achieved is very effective, because all details of the starting state, except its energy, are lost. *All* the observables, such as the pressure, the temperature, the dielectric constant, the X-ray diffraction pattern, etc., are fixed once the constraints and the energy are fixed. In the equilibrium state the number of

§§ 2.5, 19.1

thermodynamic degrees of freedom is equal to the number of constraints plus one. If we fix the constraints and one other quantity, such as the temperature, this has the effect of fixing the energy and hence the values of all other observables. Thus the nature of the equilibrium state is fixed by a relatively small number of parameters.

#### 1.4 Energy in thermodynamic systems

Modern physics has as its cornerstones the theories of quantum mechanics and relativity. The definition of energy and the fact that energy is conserved are built into the structures of these theories, and we shall not therefore spend long in discussing these two concepts here. The situation is straightforward. When forces act within a system of particles the net work done by these forces is equal to the increase in the *kinetic energy* ( $\frac{1}{2}mv^2$  summed over all particles in the non-relativistic limit). It can also be shown that, for all the forces known to physics, the work done by the forces is equal to the decrease in the corresponding *field energies*. Field energy is a form of energy which we imagine to be associated with a particular configuration of fields (such as the electric and magnetic fields) and is usually thought of as being stored in space with a definite energy density. The mutual gravitational potential energy of two massive bodies, the mutual electrostatic energy of a pair of electric charges and the energy of a light wave are all examples of field energy. Clearly, then, the sum of the kinetic energies and the field energies – the total energy – is conserved. A thermodynamic system will have inside it a definite *internal energy*,  $E$ , which is the sum of whatever particle energies and field energies are regarded as being part of the system. (We note in passing that a thermodynamic system may consist entirely of field energy: the infra-red radiation in an evacuated oven is a thermodynamic system.)

Internal energy

Scientists before the mid-nineteenth century would have been very surprised at this conclusion. They were familiar with *dissipative forces*, forces such as friction or viscous drag, which seemed to destroy or dissipate energy. If we stir a viscous liquid with a rotating paddle we certainly do *external work*,  $w$ , on it. But we also return to our starting point, so the work done cannot be stored as potential energy, apparently. Has it been lost? In a famous series of experiments Joule and others showed that in such cases *we do not in fact return to our starting point because the thermodynamic state is altered* – typically, some part of the system is hotter because all dissipative processes generate heat. Joule was able to show that a definite amount of dissipative work was always associated with a given change of state, showing that the energy was not lost, but could be accounted for in the thermodynamic change. We now know what happens. *In a dissipative process the work done is converted into the energy of disorganised small-scale motions*. The energy is not lost. It is there on an atomic scale, though it cannot be recovered in any direct way.

Dissipative work

Energy may also enter a thermodynamic system in two other ways. First, we find that when two systems are placed in *thermal contact*, energy

Heat flow

frequently flows spontaneously from one system to the other, without action by the observer. This can occur by various mechanisms, such as the transfer of vibrational energy between one solid and another whose surfaces are in contact, or the exchange of electromagnetic radiation. Such a spontaneous movement of energy is called a *heat flow*, and heat entering one body from another is written  $q$ . (The word ‘heat’ is simply a name for internal energy when we are thinking of the energy as randomised and free to move from one part of a system to another or between systems. The use of this special term can be rather misleading. There is no special part of the internal energy which can be identified as the ‘heat’, and it must not be thought that energies entering as work and as heat are stored differently inside the system.) Steps may be taken to prevent heat flow, by putting a body into a vacuum flask, for instance. Such a body is said to be *thermally isolated*. We shall refer to changes which occur inside thermally isolated systems as *adiabatic processes*. (Note carefully that authors differ over their use of this term, some reserving it for *reversible* processes in thermally isolated systems. We shall retain the earlier, more general sense.)

Adiabatic processes

Reversible work and heat flow in a slow, frictionless change of constraint

We note finally that, because the system usually exerts a force on any constraint, work may be done and energy may enter the system when a constraint is altered. For instance, a gas responds to the volume constraint by exerting a pressure  $P$ , and, if the volume changes, the work  $dw$  done on the gas by the walls of the container is equal to  $-P dV$ . There is an important difference between this type of work and the dissipative work discussed above. Dissipative processes such as stirring a viscous liquid or driving an electric current through a resistor are *irreversible*: once the work is done there is no way of recovering it by reversing the process. But imagine that we take a system, either alone or in thermal contact with a larger system which acts as a heat reservoir, and make a small and very slow frictionless change of constraint, such as an increase of volume, which we subsequently reverse. Since there is no friction we do no net work in this process, for the pressure remains at all times equal to the equilibrium pressure. Thus at the end of the process the constraints and the energy of the system (or system plus reservoir) return to their original values, and the state of the system is unchanged. The original slow and frictionless expansion is therefore *reversible*. We shall see later that if the gas is in contact with a heat reservoir during the expansion, heat will enter it from the reservoir, and leave it again if the gas is subsequently compressed. Evidently such heat flows which occur during a slow and frictionless change of constraint are also reversible.

§ 1.3

§ 5.6

To summarise, then, energy may enter a thermodynamic system in three distinct ways:

- (i) dissipative work, without change of constraint,
- (ii) heat flow, and
- (iii) reversible work, by slow, frictionless change of constraint.

Note that a *fast* movement of a constraint is not usually fully reversible, because the force involved may not be equal to the equilibrium force: in a sudden compression, for instance, the pressure may be greater than the equilibrium pressure. Such a process is best regarded as a combination of reversible and dissipative work. Evidently the change of internal energy of a system may always be written as

$$\blacktriangleright \quad dE = dw + dq, \quad (1.1) \quad \text{First law of thermodynamics}$$

where  $dw$  is the total work done on the system, either by reversible movement of the constraints or as dissipative work, and  $dq$  means the heat entering spontaneously from other systems. In thermodynamics, the principle of conservation of energy is known as the *first law*. Clearly, the first law can be expressed by Equation (1.1) combined with the *principle of conservation of heat*: heat leaving one body must reappear as heat gained by some other body or bodies.

Conservation of heat

### 1.5 Empirical temperature

Temperature is the thermodynamic parameter of which we have most intuitive knowledge: we are all aware of the sensations of hotness and coldness. But what does it mean to say that *A* is hotter than *B*? Certainly, temperature is connected with heat content and, generally speaking, adding heat to a system makes it hotter. But there are exceptions to this rule (adding latent heat to melt ice does not change the temperature), and the temperature change is not connected in any simple way with the amount of heat added. The real significance of temperature is that *it determines the direction of heat flow*. It is a remarkable fact that all equilibrium states of all systems can be placed in a sequence from 'cold' to 'hot' such that, when any two are placed in thermal contact (so that heat can pass between them) the heat *always* flows from the hotter body to the colder body, and *never* in the opposite direction. This is, of course, the basis of our physiological sense of hot and cold. A *hot* body is one from which, when we touch it, heat flows into us; a *cold* body is one which extracts heat from us. The temperature sensors in the skin respond to the change in heat density in the skin itself. We also make use of the same principle when we use a thermometer. When a thermometer is placed in contact with a warmer body, heat flows out of the body into the thermometer. In general, this makes the thermometer warmer and the body colder until they reach the same *temperature* – the same place in the sequence from cold to hot. We then (when we adopt an empirical temperature scale such as the mercury-in-glass or the perfect gas scale) use some visible parameter (such as the length of the mercury thread or the pressure of the constant volume gas thermometer) simply as a convenient label for the place on the sequence from cold to hot which the thermometer itself has reached when it is in the state concerned.

*T* fixes the direction of heat flow

Experience has shown physicists that the unwillingness of heat to flow from a colder to a hotter body applies not only to direct transfers of heat but



### 1.5 *The tasks of a mechanistic theory*

9

Second law of  
thermodynamics

also to the most elaborate indirect processes, and this conclusion is embodied in the *second law of thermodynamics*:

No process exists in which heat is transferred from a colder to a hotter body and the rest of the world is left unchanged at the end of the process.

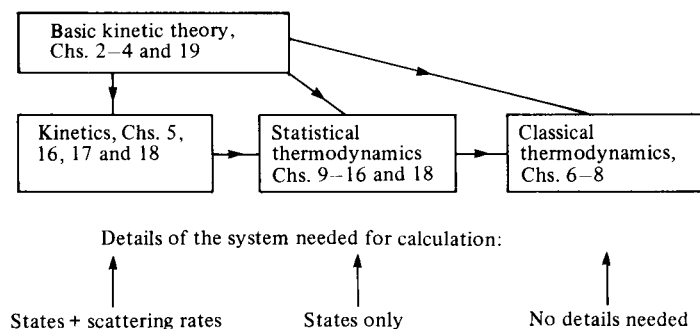
The final condition is important. In a refrigerator, for instance, heat is made to pass from a colder to a hotter body at the expense of doing work, which also appears as heat in the hotter body.

### 1.6 **The tasks of a mechanistic theory**

Some of the tasks of a mechanistic theory of heat will by now be apparent. We shall need a formalism suitable for describing the statistical state of any thermodynamic system. The formalism will have to show how, starting from an arbitrary state, systems move towards thermal equilibrium, and how fast they do so. It will have to be able, for a given energy, to predict the mean values and distributions of all possible thermodynamic observables in the equilibrium state, and to show why this state is unique. It will have to provide some means for defining and calculating from first principles the temperature of an equilibrium state, and explain the extreme reluctance of heat to flow from colder to hotter bodies. We should also expect it to account for the magnitude and frequency spectrum of thermal fluctuations.

It may be useful to mention here that the theory divides itself into three parts (Fig. 1.2). The question of how the statistical state changes with time is the subject of *kinetics*. This is the most detailed branch of thermodynamics, and involves a knowledge not only of the possible states of the system but also the mechanism by which it moves from one state to another. It is a large subject, but we shall touch on its applications only lightly in this book, in Chapters 2, 5, 16, 17 and 18. The question of the nature of the equilibrium state and the distribution of fluctuations within it is the subject usually labelled *statistical thermodynamics*. It depends on kinetics (in the sense

Fig. 1.2. The logical structure of thermodynamics, and where to find its parts in this book.



10      *Introduction: thermodynamic systems seen from outside*      1.6

that only kinetics can show which states are in equilibrium) but it turns out not to involve the details of how the system gets from one state to another, and this makes it rather more general and less system-dependent than kinetics. Certain general results of statistical thermodynamics do not depend on the details of the system concerned at all. These general results constitute the subject known as *classical thermodynamics*, traditionally derived empirically, but in this book deduced from the concepts of statistical thermodynamics. Classical thermodynamic concepts and calculations appear in Chapters 6, 7 and 8 and are used also in later chapters. The statistical foundations and the fundamental concepts of temperature and entropy on which the whole subject depends are developed in Chapters 2, 3 and 4 and briefly reassessed in the final chapter.