1

The problem

We shall be concerned in this book with the modeling of physical systems containing a large number of interacting particles. These physical systems include volumes of gases, liquids, and solids, which may or may not be isolated from their surroundings. The constituents making up these systems can be molecules, atoms, or elementary particles like electrons and protons. The techniques we will develop in this book, however, will be applicable to much broader classes of problems than those which serve as our initial focus.

1.1 A molecular model

Let us begin with a concrete physical system containing N non-relativistic, identical particles which interact pairwise. The Hamiltonian for such a system is

$$H = \frac{1}{2m} \sum_i |\mathbf{p}_i|^2 + V(\mathbf{x}_1, \ldots, \mathbf{x}_N),$$

where $\mathbf{x}_i$ and $\mathbf{p}_i$ are the positions and momenta of the particles in the system, and $i = 1, \ldots, N$. The classical approach to solving this system involves integrating Hamilton's equations of motion. A full solution, either analytic or numerical, requires that the positions and momenta of all N particles be specified at some particular time t_0. The solution must then solve the equations to trace out the full trajectories followed by all N particles in the system for some finite time interval beyond t_0.

To grasp the scope of the problem when addressed in this way, let us consider some of the numbers involved. One gram of O_2 gas contains about 2×10^{22} molecules. At typical atmospheric pressure and temperature this amount of gas will occupy a volume of about 10 cm^3. Clearly this represents a reasonably small amount of gas. We see immediately that any analysis of this system which depends on integrating the equations of motion is doomed to fail. There are just too many molecules involved. Modern computers can execute individual calculations at a rate of approximately one per nanosecond. Solving equations of motion takes many operations per molecule. But as a lower bound, suppose it took a computer a nanosecond to plot the motion of a single molecule in this gas for some period of its evolution. The same computer would then require over 300 000 years to calculate the motion of all 10^{22} molecules over the same time period!

The statistical approach to modeling systems with large numbers of degrees of freedom involves describing the system's properties using averages from a small set of appropriate and relevant observables. The thought behind the statistical approach is that the exact details of the behavior of any given degree of freedom has no significant observable effect since there are so many degrees of freedom involved. The only important variables of interest must involve averaging over many of the degrees of freedom. Statistical mechanics is the formalization of this intuitive concept. The problems to be addressed in describing statistical mechanics are threefold: under what circumstances can the properties of a physical system be defined by the behavior of an appropriate small set of variables, what are the appropriate sets of relevant variables, and how can one calculate the properties of the system in terms of these variables.

In what is to follow we will answer these questions in considerable detail. In the process we will furthermore discover a number of so-called emergent phenomena, that is physical properties, such as phase transitions or superfluidity of certain N-particle systems which are not inherent in the microscopic 1-particle Hamiltonian, but emerge in the limit when the number of particles and the size of the systems tends to infinity. The framework for our answers will be built initially from the body of work of Maxwell, Boltzmann, and Gibbs. The formal treatment of this framework begins in Chapter 2. Before beginning that formal treatment, it is necessary to describe the physical concept which is the precursor to statistical mechanics, that is classical thermodynamics. It was developed early in the nineteenth century, and is based wholly on experimental observations. What is particularly impressive about classical thermodynamics is that most of the significant results in it were developed independently of any detailed picture of the atomic nature of matter. Classical thermodynamics axiomatizes the modeling of physical systems in four basic laws. The following sections will present these laws of thermodynamics, provide modern qualitative justifications for them, and discuss some of their simple consequences.

1.2 The basics of classical thermodynamics

Consider a gas of N identical molecules contained in a volume V. The modern view of this gas is that the molecules are all in motion. For simplicity in the following discussion we will assume that the molecules all behave like hard spheres, so that molecules move in straight lines with constant velocity except during the process of collision. This motion results in constant collisions between molecules themselves, and also between the molecules and the walls which contain them in the volume V. At any given instant, one would expect that molecules will be scattered throughout the volume. Different molecules will be moving in different directions. Some molecules will be moving fast, and some will be moving slowly.

Suppose, at some particular instant, we focus on a molecule which is moving very much faster than any of its neighbors, and follow the path of this molecule for a short time. This molecule will collide with its neighbors as it moves, and because its neighbors are moving more slowly, it will tend to slow down and lose energy in these collisions. Its neighbors in turn will tend to speed up and gain energy. After a number of collisions, therefore, we can expect that our molecule will have slowed sufficiently that its momentum will become comparable to that of its neighbors, and it will no longer stand out as a fast-moving molecule.

If we focus instead on a molecule which is moving slowly relative to its neighbors then the opposite effect will occur. In collisions with its faster-moving neighbors, the slowly moving molecule will tend to pick up energy. After a number of collisions the molecule will have speeded up sufficiently that it also no longer appears exceptional when compared with its neighbors.

Thought experiments of this kind can be applied to more general situations also. Suppose for example that molecules in one half of the volume V are all moving relatively fast, while molecules in the other half of the volume are all moving relatively slowly. At the interface between the two halves of V, fast-moving molecules will be colliding with slow-moving molecules. The arguments of the previous paragraphs apply here, and we would expect that the slow-moving molecules at the interface will speed up as a result of these collisions while the fast-moving molecules will slow down. After a sufficient number of collisions, the distinction between fast- and slow-moving molecules will wash out, and the interface region will contain molecules which are moving at some common speed which is faster than the original slow-moving half volume and slower than the original fast-moving half volume. In effect, a third region will have been introduced into our system, and the original interface will have become two interfaces. The process can of course be expected to continue. After many more collisions we would expect this interface region to expand to fill the whole volume, and the original slow- and fast-moving regions to disappear and be replaced by a single region where all molecules are moving at the same typical velocity.

As long as sufficient time is allowed to elapse, we can expect that any irregularities in our system will tend to average out as a result of collisions. Individual molecules moving with large momentum relative to their neighbors will slow down. Individual molecules moving with small momentum relative to their neighbors will

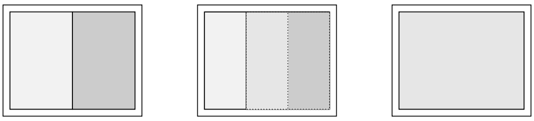

Figure 1.1 Sketch of averaging out of the density of molecules as a function of time once the separation between the two volumes is removed.

tend to speed up. Similarly, blocks of molecules with large average momentum relative to neighboring blocks will slow, losing average momentum, and blocks moving with small average momentum relative to neighboring blocks will speed up, gaining average momentum. After this sufficient time has elapsed, therefore, we can expect that the system will appear homogeneous. Molecular momenta at any one location in the system will be, on average, the same as molecular momenta at any other location.

Momentum is only one of the properties which we can expect to average out over a long timescale. Other properties such as molecular densities should average also. If the system starts with a high density of molecules in one region and a low density in a second region, then after some time we can expect that molecules will migrate out of the high-density region into the low-density region until density also averages out.

After a sufficient length of time, we can expect therefore that our system will achieve a uniform state where the only important quantities of physical interest are just those quantities which describe the typical average values in the system. For a gas of molecules, these quantities should include the density of molecules at a location, and the average momenta of those molecules.

All of the arguments just presented presuppose that the system we are describing is left undisturbed for sufficient time so that this averaging process can complete. We have implicitly presumed that there is no external source of disturbance which, for example, continuously causes molecules at some points to speed up or slow down relative to their neighbors. One of the basic assumptions of classical thermodynamics is that it is possible to arrange the environment of a physical system is such a way that no such disturbances act on it. The system is then said to be *thermally isolated*. Our thought experiments imply that, as long as sufficient time is allowed to elapse, a thermally isolated system will rearrange itself in such a way as to remove all atypical dynamics. At this time the system is said to be in *thermal equilibrium*. Thermal equilibrium is the first fundamental concept in classical thermodynamics and can be defined as follows:

Definition 1.1 A system is in *thermal equilibrium* if the state of the system does not change under time-independent external conditions.

Cambridge University Press
978-0-521-84198-6 - Elements of Statistical Mechanics: With an Introduction to Quantum Field Theory and Numerical Simulation
Ivo Sachs, Siddhartha Sen and James Sexton
Excerpt
More information

A consequence of our picture of thermal equilibrium is that systems in thermal equilibrium exhibit a uniform dynamics throughout. Since the dynamics is uniform, the important parameters needed to describe that dynamics can be encapsulated in a small number of variables, which directly or perhaps indirectly define things like the average velocity and density of particles in the system. Classical thermodynamics proposes that the different thermal equilibrium configurations of a system are distinct *thermodynamic states*. Each such state is then describable in terms of the values of a small number of *state variables*. These state variables (as we shall see later) are related to quantities such as average molecular momenta and densities. A *change of state* or *thermodynamic transformation* of a system then refers to a change in the system from one state of thermal equilibrium to a different state of thermal equilibrium. A change of state will always be signaled by a change in the value of one or more of the state variables which define the different equilibrium states.

The classical state variables which have been found experimentally to be sufficient for describing the equilibrium thermal properties of a gas of molecules are:

(1) The number N of molecules.
(2) The volume V in which the molecules are contained.
(3) The pressure P of the gas.

Pressure measures the force per unit area that the gas exerts on the surface of any object placed in the gas. It is determined by the rate at which molecules cross uniform area, and is related to the mass and average velocity of the gas molecules. In Système International (SI) units, the unit of volume is m^3. The unit of pressure is $\mathrm{N/m}^2$ or equivalently $\mathrm{kg/(m\,s^2)}$. Atmospheric pressure is of the order 10^5 $\mathrm{kg/(m\,s^2)}$ at sea level. Atmospheric pressure is usually quoted in bars where 1 bar equals 10^5 $\mathrm{kg/(m\,s^2)}$.

The state variables defining an equilibrium state are often called *macroscopic* variables of the system, since they give general information about the very large number of molecules which make a complete system. In contrast, variables which describe details of an individual molecule's dynamics, are called *microscopic* variables. Classical thermodynamics deals only in macroscopic variables, and almost exclusively with systems which are in thermal equilibrium, since state variables are only defined for such systems. A state in thermal equilibrium can be uniquely defined by specifying values for all of its state variables. What makes thermodynamics such a powerful tool is that the state variables do not depend on the history of how the system reached a given equilibrium state. Consequently one can choose a path ingeniously and in this way come to conclusions which would have been very difficult or even impossible to reach by following another path.

1.3 The zeroth law of thermodynamics

Consider now what can happen when two thermally isolated systems which are separately in equilibrium are brought into contact with each other. At the point or surface of contact, we expect that the molecules from both systems will interact with each other. These molecules may be completely free moving, or they may be bound into separate solid walls for the materials contained in each separate system. In either case, before contact, the systems are each separately in equilibrium, and the molecules at the boundaries of each system will be moving with momenta which are typical of their systems of origin. Once contact is established between the two systems, molecules from one system will start to collide with molecules from the second system. The averaging process which we have described as occurring within a system will now begin to occur between systems. The two systems start behaving like a single larger system, and so long as no further outside disturbances are added we can expect that the larger combined system will eventually smooth out its irregularities as before, and a new different thermal equilibrium will be reached for the combined system. *Thermal contact* is defined to be contact of systems, which allows this re-averaging process to occur.

In general, our intuition suggests that when two thermally isolated systems are brought into thermal contact as just described, an averaging process will result, and quantities such as average molecular momenta in each system will change as our averaging process proceeds. There is however a special case to consider, that occurs when two systems are brought into thermal contact, and no changes in the individual systems are observed. For this to happen, the average molecular velocities in the separate systems must already be correctly matched so that no change occurs when they are brought into thermal contact. This view enables us to generalize the concept of thermal equilibrium to include systems which are not originally in thermal contact with each other. Two systems which are thermally isolated and separately in thermal equilibrium are also in thermal equilibrium with each other if, when the systems are brought into thermal contact, no changes occur in the state variables describing the separate systems.

Having considered the possibilities which occur with two thermally isolated systems, the next step is to look at what happens with three or more thermally isolated systems. This leads to the zeroth law of classical thermodynamics:

The zeroth law Thermal equilibrium is transitive, i.e. if X, Y, and Z are three thermal systems, and if X is in thermal equilibrium with Y, and X is in equilibrium with Z, then Y is in thermal equilibrium with Z.

Within our molecular model, this law describes exactly what we would expect. Thermal equilibrium implies matching of quantities such as average molecular momenta between systems. If these are matched between X and Y, and between X

Cambridge University Press
978-0-521-84198-6 - Elements of Statistical Mechanics: With an Introduction to Quantum Field Theory and Numerical Simulation
Ivo Sachs, Siddhartha Sen and James Sexton
Excerpt
More information

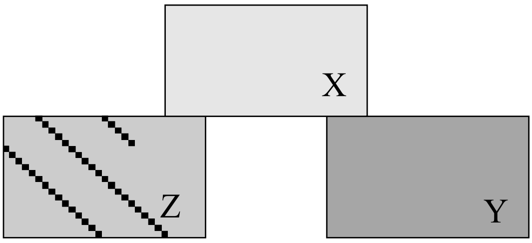

Figure 1.2 Transitivity of thermal equilibrium: If X and Y are in thermal equilibrium and X and Z are in thermal equilibrium, then Y and Z are in equilibrium.

and Z, then they must be matched between Y and Z also. As we shall immediately see, however, this law has a very profound consequence. In particular, it allows us to define a thermometer.

Systems in equilibrium, as we have seen, are described by just a few state variables. In our example these are pressure, volume, and number of particles. Different possible equilibrium configurations have different values for these state variables. So long as we only look at a single system these state variables are sufficient. If, however, we start to bring different systems into thermal contact, things get complicated. Thermal contact can change the state of a system, changing therefore the state variables describing that state. One of the fundamental pieces of information which we shall need to know is how state variables will change when systems are brought into contact. In principle, we can imagine doing a large number of experiments for all possible pairs of systems and studying what actually happens, for example, when a gas with one set of state variables is brought into thermal contact with a second gas with a different set of state variables. In practice, this is impossible. There are simply too many different possible systems, and too many ways for them to interact. The zeroth law comes to the rescue. Whenever we need to find how state variables change in a system in thermal contact, it is enough to consider the special case of interactions of the system with a special reference system. This reference system is the system X of the zeroth law and is normally called a *thermometer*.

One simple thermometer is a system where a column of mercury is confined in a long thin transparent tube. When this thermometer is brought into thermal contact with any system, the significant effect in the thermometer is that the height of the column of mercury changes until thermal equilibrium is established. This allows us to label different equilibrium states of a system by specifying the particular value of the height of the column of mercury for which equilibrium is established between the thermometer and the system. The mercury thermometer is only one of the possible thermometers that can be used. So directly indicating the particular equilibrium state by specifying this height is not reasonable. Instead we can introduce a new state variable, the *temperature*.

Definition 1.2 Two systems which are in thermal equilibrium with each other are said to have the same *temperature*.

Since temperature as defined is an abstract quantity, we are free to define its numeric value any way we wish. Some common temperature scales are defined by specifying a temperature value for water at its freezing point, and a temperature value for water at its boiling point. The Fahrenheit scale specifies the freezing point to be 32 °F, and the boiling point to be 212 °F. The Celsius scale specifies the freezing point to be 0 °C, and the boiling point to be 100 °C. Notice that the units used are completely arbitrary. We make an arbitrary decision that a particular system has a particular temperature when in one state, and a different temperature when in a different state. Other possible temperature values can then be defined by interpolating or extrapolating between these two reference values. For example a temperature of 50 °C can be defined as the temperature for which the column of mercury in a mercury thermometer is exactly halfway between its height at 0 °C and 100 °C. A thermodynamic system will have temperature 50 °C if a mercury thermometer in thermal contact with it has a column exactly this height. Clearly there is some ambiguity in these definitions. Happily, we will resolve this later in this chapter. The SI standard temperature scale is the Kelvin scale. We will define this scale explicitly later in this chapter. For the moment, we note that the Kelvin scale corresponds exactly to the Celsius scale except for an offset which is given by $0\,\mathrm{K} = -273.15\,^\circ\mathrm{C}$.

The zeroth law has allowed us to introduce a new state variable, the temperature. This variable is normally labeled T. Equilibrium states of our molecular gas are now described by four variables: N, V, P, and T. As we saw in the last section, N, V, and P are sufficient to describe the equilibrium state of a gas in isolation. We have now introduced a new state variable T which also specifies the equilibrium state. We thus expect T to be determined by the other state variables. The relationship which so determines T is called an *equation of state* and can be expressed in the generic form

$$f(T, V, P, N) = 0.$$

In particular, any substance with an equation of state of this type defines a thermometer.

For a sufficiently dilute gas at high temperature, the equation of state takes the explicit form

$$PV = NkT$$

where k is a constant of proportionality known as the Boltzmann constant. This form of the equation of state is known as the *perfect gas law*. For macroscopic systems, N is a very large number. The standard macroscopic unit of measure for molecule count is Avogadro's Number N_A given by

$$N_{\mathrm{A}} = 6.023 \times 10^{23}.$$

A mole of gas is defined to be that quantity of gas which contains N_A molecules. This quantity of gas will have mass of w grams, where w is the weight of an individual molecule expressed in atomic mass units. Using N_A we can rewrite the equation of state for a gas as

$$PV = nRT\,,$$

where $R = N_{\mathrm{A}}k$ is known as the gas constant, and $n = N/N_{\mathrm{A}}$ is the number of gram moles of the gas. The numerical values of k and R are

$$k = 1.38 \times 10^{-23}\,\mathrm{J/K} \qquad R = 8.314\,\mathrm{J/K\ per\ mol}$$

1.4 The first law, heat

Let us return now to consider what happens when two systems are brought into thermal contact. If they are in thermal equilibrium, that is, if they have the same temperature, then we know that no change of state will occur. If they are not initially in equilibrium then we expect both to change state. Since different equilibrium states are labeled by different temperatures, this change of state is usually (but not always!) indicated by a change in temperature. In this latter case, when thermal contact is established, the individual molecules will find that there is a momentum mismatch between them. Molecules in one system typically are moving slower than molecules in the second system. As the thermal contact continues, this mismatch begins to even out. Molecules in one system speed up on average, while molecules in the second system slow down on average. The immediate consequence, of course, is that a transfer of energy occurs between the two systems. Energy for a molecule in our simple example is just its momentum squared divided by twice its mass. If a molecule speeds up it gains energy, if it slows down it loses energy. Equally if all the molecules on average speed up, the energy of the system containing them increases, while if all the molecules on average slow down, the energy of the system containing them decreases.

The energy transfer which proceeds by molecular interactions of this kind is called *heat*. The symbol most often used to indicate heat in classical thermodynamics is Q. Adding heat to a system will change its internal energy, and will also change the equilibrium state of the system. As a result, when heat is added, at least some of the state variables describing the system will change. The generic case is at least that the temperature will change, and again generically this change in temperature

Cambridge University Press
978-0-521-84198-6 - Elements of Statistical Mechanics: With an Introduction to Quantum Field Theory and Numerical Simulation
Ivo Sachs, Siddhartha Sen and James Sexton
Excerpt
More information

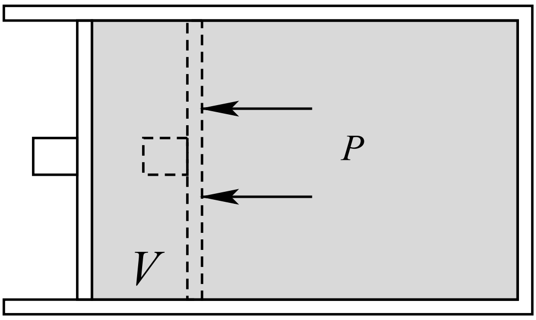

Figure 1.3 Executing work on a gas of pressure P by reducing its volume by ΔV.

is found to be proportional to the amount of heat. If ΔT is the change in temperature caused by the addition of an amount ΔQ of heat to a system, then we have

$$\Delta Q = c\Delta T.$$

The positive constant c here is the constant of proportionality and is called the *specific heat capacity* or simply *specific heat*. In general the value of c depends on the conditions under which the amount of heat ΔQ is added to the system. For gaseous systems, there are two special cases which are encountered. If ΔQ is added while keeping the volume fixed, then the specific heat is called *specific heat at constant volume*, and labeled c_V. If ΔQ is added while keeping the pressure fixed, then the the specific heat is called *specific heat at constant pressure*, and labeled c_P. One point to note is that these relations between temperature change and heat represent the generic case only. There are special cases where adding heat does not change the temperature or, even more startling, causes a negative change in temperature.

The classical unit used to measure heat is the *calorie*. This unit is defined to be the quantity of heat which must be added to one gram of water at 3.5 °C to raise its temperature by 1 °C. Heat is however a form of energy, so the modern SI unit to use for heat is the joule ($\mathrm{J} \equiv \mathrm{kg\ m^2/s^2}$). One calorie is equal to 4.1855 J.

Heat is only one way to add energy to a thermodynamic system. An alternative way is to do work on the system. For a gaseous system one simple way to do work is to force it to contract in volume. This is illustrated in Figure 1.3. Since the gas is at pressure P, it exerts a force PA perpendicular to any boundary surface of area A which contains it. If the area of the wall which is moved in this figure is A, then the distance the wall moves is $\Delta V/A$. The work done, W, by the gas is therefore force times distance moved

$$\Delta W = (PA)(\Delta V/A) = P\Delta V\,.$$