Contents

1

From Microscopic to Macroscopic Behavior

We discuss the fundamental differences between microscopic and macroscopic systems and discuss how the behavior of macroscopic systems is related to the behavior of their microscopic constituents. We note that bouncing balls come to rest and hot objects cool, and we use computer simulations to explore the qualitative behavior of macroscopic systems.

1.1 Introduction

Macroscopic systems are systems of many electrons, atoms, molecules, photons, or other constituents. Examples of familiar macroscopic objects include systems such as the air in a room, a glass of water, a metal rod, and a rubber band—examples of a gas, liquid, solid, and polymer, respectively. Less familiar macroscopic systems include superconductors, cell membranes, the brain, the economy, and neutron stars.

The questions we ask about macroscopic systems differ in important ways from the questions we ask about microscopic systems. For example, consider the air in your room. Have you ever wondered about the trajectory of a particular molecule in the air? Would knowing its trajectory be helpful in understanding the properties of air? Examples of questions that we ask about macroscopic systems include the following:

- How does the pressure of a gas depend on its temperature and the volume of its container?
- How does a refrigerator work? How can we make it more efficient?
- How much energy do we need to add to a kettle of water to change it to steam?

These questions are concerned with macroscopic properties, which describe the entire system, such as pressure, volume, and temperature, and processes related to heating and work. These questions are relevant in *thermodynamics*, which provides a framework for finding relations between the macroscopic properties of a system. Thermodynamics is concerned only with macroscopic quantities and ignores the microscopic variables that characterize individual molecules. For example, we will find in Chapter 2 that understanding the maximum efficiency of a refrigerator does not require a knowledge of the particular liquid used as the coolant. Many of the applications of thermodynamics are to engines, for example, the internal combustion engine and the steam turbine.

Another set of questions relates to understanding the behavior of macroscopic systems starting from the atomic nature of matter:

- Why do the properties of water differ from those of steam, even though water and steam consist of the same type of molecules?
- How and why does a liquid freeze into a particular crystalline structure?
- How do the properties of a system emerge from its constituents?

The laws of classical and quantum mechanics determine the behavior of molecules at the microscopic level. The goal of *statistical mechanics* is to begin with the microscopic laws of physics that govern the behavior of the constituents of the system and to deduce the properties of the system as a whole. Statistical mechanics, which we discuss in Chapters 4–9, is a bridge between the microscopic and macroscopic worlds.

Another interesting set of questions about the behavior of macroscopic systems includes the following:

- How fast does the current in a river have to be before the flow changes from laminar to turbulent?
- What will the weather be tomorrow?
- What will the climate be in a decade?

These questions are also about macroscopic phenomena, but these questions involve phenomena that are time dependent and more difficult to understand. Although there has been progress in our understanding of turbulent flow and hurricanes, it is much less advanced than our understanding of time-independent systems. Also, note the different time scales involved in understanding weather compared to climate. For this reason we will focus our attention on systems whose macroscopic properties are independent of time. However, in Chapter 10, we will consider several properties of systems in equilibrium or near equilibrium for which the time is relevant.

1.2 Some Qualitative Observations

We begin by considering a glass of hot water. We know that, if we place the glass into a large cold room, the hot water cools until its temperature equals that of the room. This simple observation illustrates two important properties of macroscopic systems—the importance of *temperature* and the *arrow of time*. Temperature is familiar, because it is associated with the physiological sensations of hot and cold and is important in our everyday experience.

The existence of the direction or arrow of time raises many questions. Have you ever observed a glass of water at room temperature spontaneously become hotter? Why not? Is there a direction of time for a single particle? Newton's second law for a single particle, $\mathbf{F} = d\mathbf{p}/dt$, implies that the motion of particles is *time-reversal invariant*; that is, Newton's second law looks the same if the time t is replaced by $-t$ and the momentum $\mathbf{p}$ is replaced by $-\mathbf{p}$. There is no direction of time at the microscopic level. Yet if we drop a basketball onto a floor, we know that it will bounce and eventually come to rest. Have you ever observed a ball at rest spontaneously begin to bounce, and then bounce higher and higher? So based on our everyday observations, we conclude that the behaviors of macroscopic bodies and single particles are very different.

Unlike scientists of about a century or so ago, we know that macroscopic systems, such as a glass of water and a basketball, consist of many molecules. Although the intermolecular forces in water produce a complicated trajectory for each molecule, the observable properties of water are easy to describe. If we prepare two glasses of water under similar conditions, we know that the observable properties of the water in each glass are indistinguishable, even though the motion of the individual particles in the two glasses is very different.

If we take into account that the bouncing ball and the floor consist of molecules, then we know that the total energy of the ball and the floor is conserved as the ball bounces and eventually comes to rest. Why does the ball eventually come to rest? You might be tempted to say the cause is "friction," but friction is just a name for an effective or phenomenological force.[1] At the microscopic level, we know that the fundamental forces associated with mass, charge, and the nucleus conserve the total energy. Hence, if we include the energy of the molecules of the ball and the floor, the total energy is conserved. Conservation of energy does not explain why the inverse process, where the ball rises higher and higher with each bounce, does not occur. Such a process also would conserve the total energy. So a more fundamental explanation is that the ball comes to rest consistent with conservation of the total energy and with some other principle of physics. We will learn that this principle is associated with an increase in the *entropy* of the system. For now, entropy is just a name,[2] and it is important to understand that energy conservation is not sufficient to understand the behavior of macroscopic systems.

Let us consider the ball bouncing on the floor again. Initially, the energy of the ball is associated with the motion of its center of mass, and we say that the energy is associated with one degree of freedom. After some time, the energy becomes associated with the individual molecules near the surface of the ball and the floor, and we say that the energy is now distributed over many degrees of freedom. If we were to bounce the ball on the floor many times, the ball and the floor would each feel warm. So we can hypothesize that energy has been transferred from one degree of freedom to many degrees of freedom, while the total energy is conserved. Hence, we conclude that the entropy is a measure of how the energy is distributed.

---

[1] We will use the word *phenomenological* often. It means a description of phenomena that is not derived from first principles.

[2] As for most concepts in physics, the meaning of entropy in the context of thermodynamics and statistical mechanics differs from its meaning as used by nonscientists.

What other quantities are associated with macroscopic systems besides temperature, energy, and entropy? We are already familiar with some of these quantities. For example, we can measure the air *pressure* in a basketball and its *volume*. More complicated quantities are the *thermal conductivity* of a solid and the *viscosity* of oil. How are these macroscopic quantities related to each other and to the motion of the individual constituent molecules? The answers to questions such as these will take us through many chapters.

1.3 Doing Work and the Quality of Energy

We have observed that hot objects cool, and cool objects do not spontaneously become hot; bouncing balls come to rest, and a stationary ball does not spontaneously begin to bounce. And although the total energy is conserved in these processes, the *distribution* of energy changes in an irreversible manner. We also have concluded that a new concept, the entropy, needs to be introduced to explain the direction of change of the distribution of energy.

Now let us take a purely macroscopic viewpoint and discuss how we can arrive at a similar qualitative conclusion about the asymmetry of nature. This viewpoint was especially important historically because of the lack of a microscopic theory of matter in the nineteenth century, when the laws of thermodynamics were being developed.

Consider the conversion of stored energy into heating a house or a glass of water. The stored energy could be in the form of wood, coal, or vegetable oil, for example. We know that this conversion is easy to do using simple methods, for example, an open flame. We also know that if we rub our hands together, they will become warmer. There is no theoretical limit to the efficiency at which we can convert stored energy to energy used for heating an object.

What about the process of converting stored energy into work? Work, like many of the other concepts that we have mentioned, is difficult to define. For now let us say that doing work is equivalent to the raising of a weight. To be useful, we need to do this conversion in a controlled manner and indefinitely. A single conversion of stored energy into work, such as the explosion of dynamite, might demolish an unwanted building, but this process cannot be done repeatedly with the same materials. It is much more difficult to convert stored energy into work, and the discovery of ways to do this conversion led to the Industrial Revolution. In contrast to the primitiveness of an open flame, we have to build an *engine* to do this conversion.

Can we convert stored energy into useful work with 100% efficiency? We know that some forms of stored energy are more useful than others. For example, why do we burn natural gas and oil in power plants even though the atmosphere and the oceans are vast reservoirs of energy? Can we mitigate global climate change by extracting energy from the atmosphere to run a power plant? From the work of Kelvin, Clausius, Carnot, and others, we know that we cannot convert stored energy into work with 100% efficiency, and we must "waste" some of the energy. At this point, it is easier to understand the reason for this necessary inefficiency by microscopic arguments. For example, the energy in the gasoline of the fuel tank of an automobile is associated with many molecules. The job of the automobile engine is to transform this (potential) energy so that it is associated with only a few degrees of freedom, that is, the rolling tires and gears. It is plausible that it is inefficient to transfer energy from many degrees of freedom to only a few. In contrast, the transfer of energy from a few degrees of freedom (the firewood) to many degrees of freedom (the air in your room) is relatively easy.

The importance of entropy, the direction of time, and the inefficiency of converting stored energy into work are summarized in various statements of the *second law of thermodynamics*.[3] We will learn that the first law is a statement of conservation of energy.

Suppose that we take some firewood and use it to "heat" a sealed room. Because of energy conservation, the energy in the room plus the firewood is the same before and after the firewood has been converted to ash. Which form of the energy is more capable of doing work? You probably agree that the firewood is a more useful form of energy than the "hot air" and ash that exists after the firewood is burned. Originally the energy was stored in the form of chemical (potential) energy. Afterward the energy is mostly associated with the motion of the molecules in the air. What has changed is not the total energy, but its ability to do work. We will learn that an increase in entropy is associated with a loss of ability to do work.

1.4 Thermal Equilibrium

So far we have discussed the behavior of macroscopic systems by appealing to our everyday experience and simple observations. We now discuss some ways of *simulating* the behavior of macroscopic systems. Although we cannot simulate a macroscopic system of 10^{23} particles, we will find that even small systems of the order of 100 particles are sufficient to illustrate the qualitative behavior of macroscopic systems.

We first discuss how we can simulate a gas consisting of molecules whose internal structure can be ignored. Imagine a system of N particles in a closed container of volume V, and suppose that the container is far from the influence of external forces such as gravity. We will usually consider two-dimensional systems so that we can easily visualize the motion of the particles.

For simplicity, we assume that the motion of the particles is given by classical mechanics, and hence we need to solve Newton's second law for each particle. To compute the total force on each particle, we have to specify the nature of the interaction between the particles. We will assume that the force between any pair of particles depends only on the distance between them. This assumption is applicable to simple liquids, such as liquid argon, but not to water. We will also assume that the particles are not charged. A potential that approximates the interaction between a pair of neutral atoms or molecules is given by the Lennard-Jones potential[4]

$$u(r) = 4\epsilon\left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6}\right], \tag{1.1}$$

---

[3] It is interesting that the second law of thermodynamics was conceived before the first law of thermodynamics.

[4] This potential is named after John Lennard-Jones, 1894–1954, a theoretical chemist and physicist. The Lennard-Jones potential is most appropriate for closed-shell systems, such as Ar or Kr.

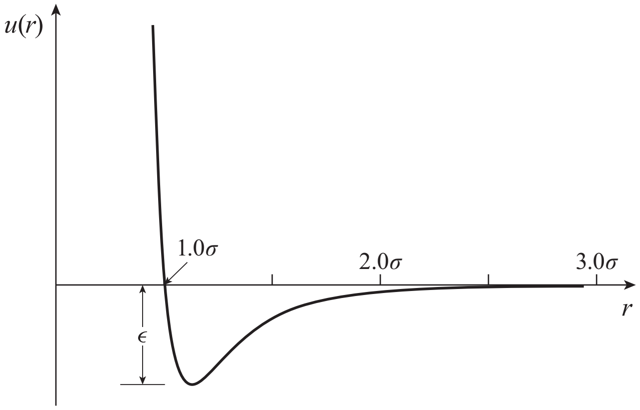

Figure 1.1. Plot of the Lennard-Jones potential $u(r)$, where r is the distance between the particles. The potential is characterized by a length σ and an energy ϵ.

where r is the distance between two particles. A plot of $u(r)$ is shown in Figure 1.1. The r^{-12} form of the repulsive part of the interaction was chosen for convenience and has no fundamental significance. The attractive $1/r^6$ behavior at large r is the van der Waals interaction.[5] The force between any two particles is given by $f(r) = -du/dr$.

In macroscopic systems, the fraction of particles near the walls of the container is negligibly small. In contrast, the number of particles that can be simulated is typically 10^3–10^6. For these small systems, the fraction of particles near the walls of the container would be significant, and the behavior of the system would be dominated by surface effects. The most common way of minimizing surface effects and to simulate more closely the properties of a macroscopic system is to use what are known as toroidal boundary conditions. These boundary conditions are familiar to computer game players. For example, a particle that exits the right edge of the "box" reenters the box from the left side. In one dimension, this boundary condition is equivalent to taking a piece of string and making it into a loop. In this way a particle moving on the string never reaches the end (the surface).

Given the form of the interparticle potential, we can determine the total force on each particle due to all the other particles in the system. We then use Newton's second law of motion to find the acceleration of each particle. Because the acceleration is the second derivative of the position, we need to solve a second-order differential equation for each particle (in each direction), given the initial position and velocity of each particle. (For a two-dimensional system of N particles, there are $2N$ differential equations.) These differential equations are coupled, because the acceleration of one particle depends on the positions of all the other particles. Although we cannot solve the coupled differential equations analytically, we can use numerical methods to solve them to a good approximation. This way of simulating gases, liquids, solids, and biomolecules is called *molecular dynamics*.

In the following, we will explore some of the qualitative properties of macroscopic systems by doing some simple simulations. Before you do the simulations, think

---

[5] The van der Waals interaction arises from an induced dipole-dipole effect. It is present in all molecules, but is important only for the heavier noble gas atoms.

about what you believe the results will be. In many cases the most valuable part of the simulation is not the simulation itself, but the act of thinking about a concrete model and its behavior.

Problem 1.1. Approach to equilibrium
Suppose that we divide a box into three equal parts and place N particles in the middle third of the box.[6] The particles are placed at random with the constraint that no two particles can be closer than the length parameter σ. This constraint prevents the initial force between any two particles from being too big, which would lead to the breakdown of the numerical method used to solve the differential equations. The velocity of each particle is assigned at random and then their velocities are shifted so that the velocity of the center of mass is zero. At $t = 0$, we remove the "walls" between the three parts and watch the particles move according to Newton's equations of motion. We say that the removal of a wall corresponds to the removal of an internal constraint. What do you think will happen?

Program `ThreePartsMD` implements this simulation. The program shows the positions of the particles and plots the number of particles in the left (n_1), center (n_2), and right (n_3) part of the box as a function of time. The input parameter is N, the number of particles initially in the center, so that $n_1 = 0$, $n_2 = N$, and $n_3 = 0$ at $t = 0$.

(a) Does the system appear to show a direction of time for $N = 6$?

(b) Does the system appear to show a direction of time for $N = 27$? Choose various values of N that are multiples of 3 up to $N = 270$. What is the nature of the time dependence of n_1, n_2, and n_3? Is the direction of time better defined for larger N?

(c) To better understand the direction of time, make a video of the motion of the positions of 270 particles starting from $t = 0$. Run the simulation until the particles are approximately equally divided between the three regions. Then run the video both forward and backward. Can you tell by looking at the video which direction is forward? Repeat for larger values of N. Does your conclusion about the direction of time become more certain for larger N?

(d) Repeat part (c) but start the video after the particles are distributed equally among the three regions, say at $t = 20$, as in Figure 1.2. Is the direction of time more apparent? Repeat for various values of N.

(e) The time shown in the plots is in terms of $\sigma(m/\epsilon)^{1/2}$, where σ and ϵ are the length and energy parameters of the Lennard-Jones potential in (1.1), and m is the mass of a particle. Verify that this combination has units of time. For argon, $\sigma = 3.4 \times 10^{-10}$ m, $\epsilon = 1.65 \times 10^{-21}$ J, and $m = 6.69 \times 10^{-26}$ kg. What is the value of $\sigma(m/\epsilon)^{1/2}$ for argon? How much real time has elapsed if the program shows that $t = 100$?

(f) Program `TwoPartsMD` initially divides the box into two parts rather than three. Run the program and verify that the simulation shows similar qualitative behavior. Explain the use of toroidal boundary conditions.

---

[6] We divided the box into three parts so that the effects of the toroidal boundary conditions are not as apparent as if we had initially confined the particles to one half of the box.

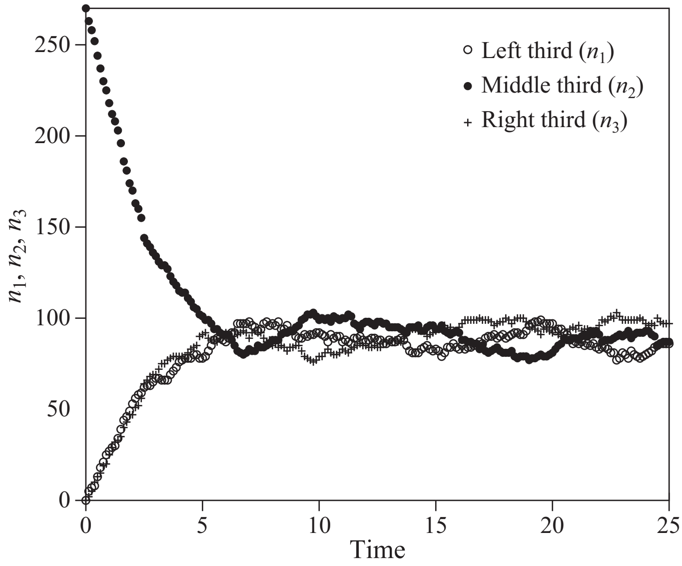

Figure 1.2. Evolution of the number of particles in each third of the box for $N=270$. The particles were initially restricted to the middle third of the box. Toroidal boundary conditions are used in both directions. The initial velocities are assigned at random such that the center of mass velocity is zero.

(g) Why does the system exhibit a direction of time when the motion of each particle is time reversible?

(h) *After n_1, n_2, and n_3 first become approximately equal for $N=270$, reverse the velocities of all the particles and continue the simulation. Reversing the velocities is equivalent to seeing the simulation go backward in time. Do the particles return to the middle third of the box? Do the simulation again, but let the system evolve longer before reversing the velocities. Are your results any different? Are your results consistent with the fact that Newton's equations of motion are deterministic and time reversal invariant, so that reversing the velocities should exactly retrace the original particle trajectories? (See Section 1.7 for a discussion of the extreme sensitivity of the trajectories to very small errors.) □

1.4.1 A probabilistic model

To gain some more insight into why there is a direction of time, we consider a simpler model that shows similar behavior. Imagine a closed box that is divided into two parts of equal areas. The left half initially contains N identical particles, and the right half is empty. We then make a small hole in the wall between the two halves. Instead of simulating this system by solving Newton's equations for each particle, we adopt an approach based on a probabilistic model. We assume that the system is so dilute that the particles do not interact with one another. Hence, the probability per unit time that a particle goes through the hole in the wall is the same for all particles regardless of the number of particles in either half. We also assume that the size of the hole is such that only one particle can pass through in one unit of time.

One way to implement this model is to choose a particle at random and move it to the other half. However, our interest is only in the number of particles on each side, not on which particles are on each side. Hence, we need to know only n, the number of particles on the left side; the number on the right side is $N-n$. Because each particle has the same chance to go through the hole in the wall, the probability per unit time that a particle moves from left to right equals the number of particles on the left side divided by the total number of particles. That is, the probability of a move from left to right is n/N. The algorithm for simulating the evolution of the model is given by the following steps:

1. Generate a random number r from a uniformly distributed set of random numbers in the interval $0 \le r < 1$.
2. If $r \le n/N$, a particle is moved from left to right, that is, let $n \to n-1$; otherwise, $n \to n+1$.
3. Increase the time by 1.

Program `HoleInWall` implements this algorithm and plots the evolution of n.

Problem 1.2. Hole in the wall

(a) Before you run the program, describe what you think will be the qualitative behavior of $n(t)$, the time dependence of the number of particles on the left side of the box.

(b) Run the program, and describe the behavior of $n(t)$ for various values of N. Does the system approach equilibrium? How would you characterize equilibrium? In what sense is equilibrium better defined as N is increased? Does your definition of equilibrium depend on how the particles were initially distributed between the two halves of the box?

(c) When the system is in equilibrium, does the number of particles on the left-hand side remain a constant? If not, how would you describe the nature of equilibrium?

(d) If $N \gtrsim 32$, does the system return to its initial state during the time you have patience to watch the system?

(e) How does $\overline{n}$, the average number of particles on the left-hand side, depend on N after the system has reached equilibrium? The program computes various averages from the time $t=0$. Why does such a calculation not yield the correct equilibrium average values? Use the `Zero Averages` button to reset the calculation of the averages.

(f) Define the quantity σ by the relation[7]

$$\sigma^2 = \overline{(\Delta n)^2} = \overline{(n-\overline{n})^2}. \tag{1.2}$$

What does σ measure? What would be its value if n were constant? How does σ depend on N? How does the ratio $\sigma/\overline{n}$ depend on N? We say that σ is a measure of the fluctuations of n about its average, and $\sigma/\overline{n}$ is a measure of the relative fluctuations of n. ☐

---

[7] This use of σ should not be confused with the length σ in the Lennard-Jones potential.

TABLE 1.1
The four possible ways in which $N=2$ particles can be distributed between the two halves of a box. The quantity $W(n)$ is the number of microstates corresponding to the macroscopic state characterized by n, the number of particles on the left-hand side. The probability $P(n)$ of macrostate n is calculated assuming that each microstate is equally likely.

Microstate		n	$W(n)$	$P(n)$
L	L	2	1	1/4
L	R	1	2	1/2
R	L			
R	R	0	1	1/4

From Problems 1.1 and 1.2 we conclude that the average values of the macroscopic quantities of interest will eventually become independent of time. We say that the system has reached equilibrium, and the macroscopic quantities exhibit fluctuations about their average values. We also learned that the relative fluctuations become smaller as the number of particles is increased, and the details of the dynamics are irrelevant to the general tendency of macroscopic systems to approach equilibrium.

How can we understand why the systems we have considered exhibit a direction of time? There are two general approaches. One approach is to study the dynamics of the system.[8] A simpler way is to change the question and take advantage of the fact that the equilibrium state of a macroscopic system is independent of time on the average, and hence time is irrelevant in equilibrium. For the model considered in Problem 1.2, we will see that counting the number of ways that the particles can be distributed between the two halves of the box will help us understand the nature of equilibrium. This approach does not tell us how long it takes the system to reach equilibrium, but it will give us insight into why there is a direction of time.

1.4.2 Counting states

We call each distinct arrangement of the particles between the two halves of the box a *microstate*. For $N=2$ the four possible microstates are shown in Table 1.1. Because the halves are equivalent, a given particle is equally likely to be in either half when the system is in equilibrium. Hence, for $N=2$ the probability of each microstate equals 1/4 when the system is in equilibrium.

From a macroscopic point of view, we do not care which particle is in which half of the box, but only the number of particles on the left. Hence, the macroscopic state or *macrostate* is specified by n. Are the three possible macrostates listed in Table 1.1 equally probable?

We next consider $N=4$, for which there are $2 \times 2 \times 2 \times 2 = 2^4 = 16$ microstates (see Table 1.2). We assume as before that all microstates are equally probable in equilibrium. We see from Table 1.2 that there is only one microstate with all particles on the left, and the most probable macrostate is $n=2$.

For larger N the probability of the most probable macrostate with $n=N/2$ is much greater than the macrostate with $n=N$, which has a probability of only

---

[8] The approach to equilibrium of the number of particles in the left half of the box in Problem 1.2 is discussed in Section 1.14.1.

TABLE 1.2
The 2^4 possible ways in which $N = 4$ particles can be distributed between the two halves of a box. The quantity $W(n)$ is the number of microstates corresponding to the macroscopic state characterized by n. The probability $P(n)$ of macrostate n is calculated assuming that each microstate is equally likely.

Microstate				n	$W(n)$	$P(n)$
L	L	L	L	4	1	1/16
R	L	L	L	3		
L	R	L	L	3		
L	L	R	L	3		
L	L	L	R	3		
					4	4/16
R	R	L	L	2		
R	L	R	L	2		
R	L	L	R	2		
L	R	R	L	2		
L	R	L	R	2		
L	L	R	R	2		
					6	6/16
R	R	R	L	1		
R	R	L	R	1		
R	L	R	R	1		
L	R	R	R	1		
					4	4/16
R	R	R	R	0	1	1/16

$1/2^N$, corresponding to a single microstate. Hence, we conclude that the equilibrium macrostate corresponds to the most probable state.

Problem 1.3. Counting microstates

(a) Calculate the number of possible microstates for each macrostate n for $N = 8$ particles. What is the probability that $n = 8$? What is the probability that $n = 4$? It is possible to count the number of microstates for each n by hand, but because there are a total of $2^8 = 256$ microstates, counting this way is tedious. An alternative is to determine the number of ways that n particles out of N can be in the left half of the box. Enumerate the possible microstates for smaller values of N until you see a pattern.

(b) The macrostate with $n = N/2$ is much more probable than the macrostate with $n = N$. Why? □

1.4.3 Some qualitative observations

Approach to equilibrium. The macrostates that give us the least amount of information about the associated microstates are the most probable. For example, suppose that we wish to know where particle 1 is, given that $N=4$. If $n=4$, we know with certainty that particle 1 is on the left. If $n=3$, the probability that particle 1 is on the left is 3/4. And if $n=2$, we know only that particle 1 is on the left with probability 1/2. In this sense the macrostate $n=2$ is more random than macrostates $n=4$ and $n=3$.

We also found from the simulations in Problems 1.1 and 1.2 that if an isolated macroscopic system changes in time due to the removal of an internal constraint, it tends to evolve from a less random to a more random state. Once the system reaches its most random state, fluctuations corresponding to an appreciably nonuniform state are rare. These observations and our reasoning based on counting the number of microstates corresponding to a particular macrostate leads us to conclude that:

> *A system in a nonuniform macrostate will change in time on the average so as to approach its most random macrostate where it is in equilibrium.*

This conclusion is independent of the nature of the dynamics. Note that the simulations in Problems 1.1 and 1.2 involved the dynamics, but our discussion of the number of microstates corresponding to each macrostate did not involve the dynamics in any way. Instead we counted (enumerated) the microstates and assigned them equal probabilities, assuming that the system is isolated and in equilibrium.

In the simulations discussed in Problem 1.1, the total energy was conserved, and hence the macroscopic quantity of interest that changed from the specially prepared initial state to the most probable macrostate was not the total energy. What macroscopic quantities changed besides the number of particles in each part of the box? Based on our previous discussion, we can tentatively say that the quantity that changed is the entropy. We conjecture that the entropy is associated with the number of microstates associated with a given macrostate. If we make this association, we conclude that the entropy is greater after the system has reached equilibrium than in the system's initial state. Moreover, if the system was initially prepared in a microstate such that $n_1=n_2=n_3=N/3$, the average values of $n_1=n_2=n_3=N/3$ and hence the entropy would not change. Hence, we can conclude the following:

> *The entropy of an isolated system increases or remains the same when an internal constraint is removed.*

This statement is equivalent to the second law of thermodynamics. This identification of the entropy is made explicit in Chapter 4.

As a result of the simulations that we have done and our discussions, we can make some additional preliminary observations about the behavior of macroscopic systems.

Fluctuations in equilibrium. Once a system reaches equilibrium, the macroscopic quantities of interest are not independent of time, but *exhibit fluctuations about their average values*. Only the *average* values of the macroscopic variables are independent of time in equilibrium.[9] For example, in Problem 1.2 the number of particles $n(t)$ changes with t, but its average value $\overline{n}$ does not. If N is large, fluctuations corresponding to a very nonuniform distribution of the particles almost never occur, and the relative fluctuations $\sigma/\overline{n}$ become smaller as N is increased.

History independence. The properties of *equilibrium systems are independent of their history*. For example, the value of $\overline{n}$ in Problem 1.2 would be the same whether we had started with $n(t=0)=N$ or $n(t=0)=0$. In contrast, we are all products of our history. One consequence of history independence is that we can ignore how a system reached equilibrium.

Need for statistical approach. A macroscopic system can be described in detail by specifying its *microstate*. Such a description corresponds to specifying all the possible information. For a system of classical particles, a microstate corresponds to specifying the position and velocity of each particle.

Suppose that we simulate a system of many particles and save the trajectories of each particle as a function of time. What could we do with this information? If the number of particles is 10^6 or more or if we ran the simulation long enough, we would have a problem storing the data. Do we want to have a detailed description of the motion of each particle? Would this data give us much insight into the macroscopic behavior of the system? We conclude that the presence of a large number of particles motivates us to adopt a statistical approach. In Section 1.7, we will discuss another reason why a statistical approach is necessary.

We will find that the laws of thermodynamics depend on the fact that the number of particles in macroscopic systems is enormous. A typical measure of this number is the Avogadro number, which is approximately 6×10^{23}, the number of atoms in a mole. When there are so many particles, predictions of the average properties of the system become meaningful, and deviations from the average behavior become less important as the number of atoms is increased.

Equal a priori probabilities. In our analysis of the probability of each macrostate in Problem 1.2, we assumed that each microstate is equally probable. That is, each microstate of an isolated system occurs with equal probability if the system is in equilibrium. We will make this assumption explicit in Chapter 4.

1.5 Measuring the Pressure and Temperature

The obvious macroscopic quantities that we can measure in our simulations of many particles include the average kinetic and potential energies. We also know the number of particles and the volume. From our everyday experience we know that there are

---

[9] Some of our general statements will be qualified in later chapters. Just because the macroscopic properties of a system are time independent does not necessarily mean that it is in equilibrium. If a system is driven by external forces or currents that are time independent, the observable macroscopic properties of the system can be time independent, and the system is said to be in a *steady state*. For example, consider a metal bar with one end in contact with a much larger system at temperature T_{hot} and the other end in contact with a much larger system at temperature T_{cold}. If $T_{\text{hot}} > T_{\text{cold}}$, energy will be continually transported from the "hot" end to the "cold" end, and the temperature gradient will not change on average. We discuss this situation in Chapter 10.

at least two other macroscopic variables that are relevant for a system of particles, namely, the pressure and the temperature.

You are probably familiar with the concepts of force and pressure from courses in mechanics. The force on a wall of a container is related to the rate of change of the component of the total momentum perpendicular to the wall. Because the force on the wall is proportional to the area A of the wall, we define the pressure P as $P = F/A$. The pressure is a scalar, because it is the same in all directions on the average.[10]

The number of particles that strike a wall of the box per second is huge. A pressure gauge cannot distinguish between the individual frequent impacts but measures the average force due to many impacts. In this book, we will discuss many examples of the relation of the macroscopic properties of a system to an average of some microscopic quantity.

Before we discuss the nature of temperature, consider the following questions.

Problem 1.4. Nature of temperature

(a) Summarize what you know about temperature. What reasons do you have for thinking that it has something to do with energy?

(b) Discuss what happens to the temperature of a hot cup of coffee. What happens, if anything, to the temperature of its surroundings?

(c) If you add energy to a pot of boiling water, does the temperature of the water change? □

Although temperature and energy are related, they are not the same quantity. For example, one way to increase the energy of a glass of water is to lift it. However, this action would not affect the temperature of the water. Also, if we placed a glass of water on a moving conveyor belt, the temperature of the water would not change. So the temperature has nothing to do with the motion of the center of mass of the system. We also know that temperature is a property associated with many particles. It would be absurd to refer to the temperature of a single molecule.

The most fundamental property of temperature is not that it has something to do with energy. More importantly, *temperature is the quantity that becomes equal when two systems are allowed to exchange energy with one another*. In Problem 1.5, we interpret the temperature from this point of view.

Problem 1.5. Identification of the temperature

(a) Use program `ThermalContact` to simulate two systems, A and B, of particles that interact via the Lennard-Jones potential in (1.1). Both systems are in a square box with linear dimension $L = 12$. Toroidal boundary conditions are not used, and the particles also interact with fixed particles (with infinite mass) that make up the walls and the barrier between them. Initially, the two systems are isolated from each other and from their surroundings. We take $N_A = 81$, $N_B = 64$, $\epsilon_{AA} = 1.0$, $\sigma_{AA} = 1.0$, $\epsilon_{BB} = 1.5$, $\sigma_{BB} = 1.2$, $\epsilon_{AB} = 1.25$, and $\sigma_{AB} = 1.1$. Run the simulation and monitor the kinetic energy and potential energy until each system appears to reach equilibrium. What is the

---

[10] Because most of our simulations are done using toroidal boundary conditions, we will use the relation of the force to the *virial*, a mechanical quantity that involves all the particles in the system, not just those colliding with a wall. See Gould, Tobochnik, and Christian, Chap. 8 or an advanced classical mechanics text.

average potential and kinetic energy of each system? Is the total energy of each system constant (to within numerical error)?

(b) Remove the barrier and let the particles in the two systems interact with one another. What quantity is exchanged between the two systems? (The volume of each system is fixed.)

(c) After equilibrium has been established, compare the average kinetic and potential energies of each system to their values before the two systems came into contact.

(d) What quantity is the same in both systems after equilibrium has been established? Are the average kinetic and potential energies the same? What quantities would change if you doubled the number of particles and the area of each system? Would the temperature change? Does it make more sense to compare the average kinetic and potential energies or the average kinetic and potential energies per particle? Do any other quantities become approximately equal? What can you conclude about the possible identification of the temperature in this system?

After the barrier has been removed, the average kinetic energy per particle for both systems becomes equal. It is natural to assume that the temperature is proportional to the average kinetic energy per particle, which is valid for systems of particles for which quantum effects can be ignored.

You might have noticed that the final temperature is slightly higher than the initial temperatures of each system. This behavior seems inconsistent with what we observe in nature. That is, when two systems are placed in thermal contact, the higher temperature decreases and the lower temperature increases. The reason for this inconsistency with what we observe in nature is that when we allow both sides of the system to interact in the simulation, we have added the extra interaction energy to the composite system. □

From the simulations in Problem 1.5 you likely concluded that the temperature is proportional to the average kinetic energy per particle. You might have known about this relation already. We will learn in Chapter 6 that the proportionality of the temperature to the average kinetic energy per particle holds only for a system of particles whose kinetic energy is proportional to the square of its momentum or velocity.

The relation between temperature and kinetic energy does not tell us why temperature is the quantity that determines the direction of energy transfer when two systems are in thermal contact. The temperature of an object is a measure of how easy it is for the object to accept more energy when placed in thermal contact. Low-temperature objects easily accept energy, and high-temperature objects easily give up energy. When two objects are at the same temperature, they do not exchange a net amount of energy.

How can we measure the temperature of a system? After all, in an experiment we cannot directly measure the average kinetic energy per particle. Nevertheless, there are many kinds of thermometers. These thermometers exchange energy with the system of interest and have some physical property that changes in a way that can be calibrated to yield the temperature. In Problem 1.6 we ask you to think about the general characteristics of thermometers. We then consider a simple model thermometer. We will discuss thermometers in more detail in Chapter 2.

Problem 1.6. Thermometers

(a) Describe some of the thermometers with which you are familiar.
(b) On what physical principles do these thermometers operate?
(c) What requirements must a thermometer have to be useful? □

To gain more insight into the meaning of temperature, we consider a model thermometer known as a "demon." The demon is a special particle that carries a sack of energy and exchanges energy with the system of interest by making a small change in the system. If the change lowers the energy of the system, the demon puts the extra energy in its sack. If the change would increase the energy of the system, the demon gives the system the needed energy with the constraint that E_d, the energy of the demon, must be nonnegative. The total energy of the demon and the system is a constant.

The behavior of the demon is given by the following algorithm:

1. Choose a particle in the system at random and make a trial change in one of its coordinates.
2. Compute ΔE, the change in the energy of the system due to the trial change.
3. If $\Delta E \leq 0$, the system gives the surplus energy to the demon, $E_d \to E_d + |\Delta E|$, and the trial change is accepted.
4. If $\Delta E > 0$ and the demon has sufficient energy (remember that $E_d \geq 0$), then the demon gives the necessary energy to the system, $E_d \to E_d - \Delta E$, and the trial change is accepted. Otherwise, the trial change is rejected and the microstate is not changed.
5. Repeat steps 1–4 many times.
6. Compute the averages of the quantities of interest once the system and the demon have reached equilibrium.

We consider the consequences of these simple rules in Problem 1.7. The nature of the demon is discussed further in Section 4.9.

Problem 1.7. The demon and the ideal gas
Program `DemonIdealGas` simulates a demon that exchanges energy with an ideal gas of N particles in d spatial dimensions. Because the particles do not interact, the only coordinate of interest is the velocity of the particles. The demon chooses a particle at random and changes each component of its velocity by an amount chosen at random between $-\Delta$ and $+\Delta$. For simplicity, we set the initial demon energy $E_d = 0$ and the initial velocity of each particle equal to $+v_0\hat{x}$, where $v_0 = (2E_0/mN)^{1/2}$. E_0 is the desired total energy of the system, and m is the mass of the particles. We choose units such that $m = 1$, and the energy and momentum are dimensionless (see Section 1.12).

(a) Run the simulation using the default parameters $N = 40$, $E_0 = 40$, and $d = 3$. Does the average energy of the demon approach a well-defined value after a sufficient number of energy exchanges with the system? (One time unit is equal to N trial changes.)
(b) What is $\overline{E}_d$, the average energy of the demon, and $\overline{E}$, the average energy of the system? Compare the values of $\overline{E}_d$ and $\overline{E}/N$.

(c) Fix $N = 40$ and double the total energy of the system. (Remember that $E_d = 0$ initially.) Compare the values of $\overline{E}_d$ and $\overline{E}/N$. Consider other values of $N \geq 40$ and E and determine the relation between $\overline{E}_d$ and $\overline{E}/N$.[11]

(d) You have probably learned in other courses that the average energy of an ideal gas in three dimensions is equal to $\overline{E} = \frac{3}{2}NkT$, where T is the temperature of the gas, N is the number of particles, and k is a constant. Our choice of dimensionless variables implies that we have chosen units such that $k = 1$. Use this relation for $\overline{E}$ to determine the temperature of the ideal gas in parts parts (b) and (c). Is $\overline{E}_d$ proportional to the temperature of the gas?

(e)* Suppose that the energy-momentum relation of the particles is not $\epsilon = p^2/2m$, but is $\epsilon = cp$, where c is a constant (which we take to be 1). Consider various values of N and E as you did in part (c). Is the dependence of $\overline{E}_d$ on $\overline{E}/N$ the same as you found in part (c)? We will show in Problem 4.30 that $\overline{E}_d$ is still proportional to the temperature.

(f)* After the demon and the system have reached equilibrium, we can compute the histogram $H(E_d)\Delta E_d$, the number of times that the demon has energy between E_d and $E_d + \Delta E_d$. The bin width ΔE_d is chosen by the program. The histogram is proportional to the probability $p(E_d)\Delta E$ that the demon has energy between E_d and $E_d + \Delta E$. What do you think is the dependence of $p(E_d)$ on E_d? Is the demon more likely to have zero or nonzero energy?

(g)* Verify the exponential form, $p(E_d) = Ae^{-\beta E_d}$, where A and β are parameters. How does the value of $1/\beta$ compare to the value of $\overline{E}_d$? We will find that the exponential form of $p(E_d)$ is universal, that is, independent of the system with which the demon exchanges energy, and that $1/\beta$ is proportional to the temperature of the system.

(h) Discuss why the demon is an ideal thermometer. □

1.6 Work, Heating, and the First Law of Thermodynamics

As you watch the motion of the individual particles in a molecular dynamics simulation, you will probably describe the motion as "random" in the sense of how we use random in everyday speech. The motion of the individual molecules in a glass of water exhibits similar motion. Suppose that we were to expose the water to a low flame. In a simulation this process would roughly correspond to increasing the speed of the particles when they hit a wall. We say that we have transferred energy to the system *incoherently*, because each particle continues to move more or less at random.

In contrast, if we squeeze a plastic container of water, we do *work* on the system, and the particles near the wall move *coherently*. So we can distinguish two different ways of transferring energy to the system. *Heating transfers energy incoherently and doing work transfers energy coherently.*

Let us again consider a molecular dynamics simulation and suppose that we have increased the energy of the system by either compressing the system from a volume V_i to V_f and doing work on it or by randomly increasing the speed of the particles

---

[11] Because there are finite-size effects that are order $1/N$, it is desirable to consider $N \gg 1$. The trade-off is that the simulation takes longer to run.

that reach the walls of the container with a volume V_f. Roughly speaking, the first way would initially increase the potential energy of interaction, and the second way would initially increase the kinetic energy of the particles. If we increase the total energy by the same amount, could you tell by looking at the particle trajectories after equilibrium has been reestablished how the energy was increased? The answer is no, because for a given total energy, volume, and number of particles, the kinetic energy and the potential energy each have unique equilibrium values. We conclude that the total energy of the gas can be changed by doing work on it or by heating it or by both processes. This statement is equivalent to the *first law of thermodynamics* and from the microscopic point of view is a statement of conservation of energy.

Our discussion implies that the phrase "adding heat" to a system makes no sense, because we cannot distinguish "heat energy" from potential energy and kinetic energy. Nevertheless, we frequently use the word "heat" in everyday speech. For example, we might say "Please turn on the heat" and "I need to heat my coffee." We will avoid such uses, and whenever possible avoid the use of "heat" as a noun. Why do we care? Because there is no such thing as heat, and the words we use affect how we think.

1.7 *The Fundamental Need for a Statistical Approach

As discussed in Section 1.4, we compute the position and velocity of every particle in a molecular dynamics simulation. Our disinterest in the trajectory of any particular particle and the overwhelming amount of information that is generated motivates us to develop a statistical approach.

We now discuss the more fundamental reason we must use probabilistic methods to describe systems with more than a few particles. The reason is that for a wide variety of conditions, even the most powerful supercomputer yields positions and velocities that are meaningless! We will find that the trajectories in a system of many particles depend sensitively on the initial conditions. Such a system is said to be *chaotic*. This behavior forces us to take a statistical approach even for some systems with as few as three particles.[12]

As an example, consider a system of $N = 11$ particles interacting with the Lennard-Jones potential (see Program `Chaos`). The initial conditions are such that all the particles have the same velocity $v_x(i) = 1$, $v_y(i) = 0$, and the particles are equally spaced vertically, with $x(i) = L/2$ for $i = 1, \ldots, 11$ [see Figure 1.3(a)]. For these special initial conditions, the particles will continue moving indefinitely in the x direction (using toroidal boundary conditions).

We now stop the simulation and change the velocity of particle 6 so that $v_x(6) = 1.000001$. What do you think will happen? In Figure 1.3(b), we show the positions of the particles at $t = 8.0$ after the change in velocity of particle 6. Note that the positions of the particles are no longer equally spaced. A small change in the velocity of one particle led to a big change in the trajectories of all the particles.

---

[12] The planetary system is not immune from chaos. For example, in 10^{12} years, Mercury may suffer a close encounter with Venus or plunge into the Sun. However, we can predict the trajectories of the planets for hundreds of thousands of years. See for example, M. Lecar, F. A. Franklin, and M. J. Holman, "Chaos in the solar system," Annu. Rev. Astron. Astrophys. **39**, 581–631 (2001).

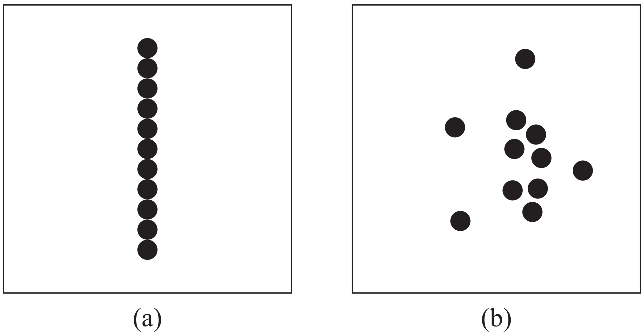

Figure 1.3. (a) A special initial condition for $N = 11$ particles such that their motion remains parallel indefinitely. (b) The positions of the particles shortly after a small change in $v_x(6)$ from 1 to 1.000001.

Problem 1.8. Irreversibility
Program `Chaos` simulates a system of $N = 11$ particles with the special initial condition shown in Figure 1.3(a).

(a) Perturb the velocity of particle 6, and confirm the qualitative results that we have discussed.
(b) Stop the simulation at a time t after the perturbation, and reverse all the velocities. Confirm that if t is sufficiently short, the particles will return approximately to their initial state. What is the maximum value of t that allows the particles to return to their initial positions if t is replaced by $-t$ (all velocities reversed)? □

An important property of chaotic systems is their *extreme sensitivity to initial conditions*, that is, the trajectories of two identical systems starting with slightly different initial conditions will diverge exponentially. For such systems we cannot predict the positions and velocities of the particles very far into the future because even the slightest error in our measurement of the initial conditions would make our prediction incorrect if the elapsed time is sufficiently long. That is, we cannot answer the question "Where is a particular particle at time t?" if t is sufficiently long. It might be disturbing to realize that our answers are meaningless if we ask the wrong questions.

Suppose that a chaotic system evolves for a time t and all the velocities are reversed. If the system is allowed to evolve for an additional time t, the system will not return to its original state unless the velocities are specified with infinite precision. This lack of practical reversibility is related to what we observe in macroscopic systems. If you pour milk into a cup of coffee, the milk becomes uniformly distributed throughout the cup. You will never see a cup of coffee spontaneously return to the state where all the milk is at the surface, because the positions and velocities of the milk and coffee molecules would need to have exactly the right values to allow the milk to return to this very special state. This sensitivity to initial conditions provides the foundation for the arrow of time.

1.8 *Time and Ensemble Averages

Because of the sensitivity to initial conditions of the dynamics of many-particle systems, the computed trajectories will not be exact. However, the computed trajectories can be made consistent with the constraints of constant total energy E, volume V, and number of particles N. Because of the assumption of equal probabilities for the microstates of an isolated system, all trajectories consistent with the constraints will contribute equally to the average values of macroscopic quantities. Thus, the trajectories computed in a simulation are meaningful, even though they are very different from the exact trajectories that would be calculated with an infinite-precision computer.

Molecular dynamics simulations allow us to calculate a time average. If we do a laboratory experiment to measure the temperature and pressure, our results also would be equivalent to a time average. Because time is irrelevant for a system in equilibrium, we can also calculate an *ensemble* average of various quantities. An ensemble average is calculated at a given time by considering many identical systems that satisfy the same known conditions.

An example might clarify the nature of the two types of averages. Suppose that we want to determine the probability that the toss of a coin results in "heads." We can do a time average by taking one coin, tossing it in the air many times, and counting the fraction of heads. In contrast, an ensemble average can be found by obtaining many similar coins and tossing them into the air at one time. We will discuss ensemble averages in Chapter 3.

It is reasonable to assume that the two types of averages are equivalent. This equivalence is called the *ergodic hypothesis*. The term "hypothesis" might suggest that the equivalence is not well accepted, but it reminds us that the equivalence has been shown to be rigorously true in only a few cases. The sensitivity of the trajectories of chaotic systems to initial conditions suggests that a classical system of particles moving according to Newton's equations of motion passes through many different microstates corresponding to different sets of positions and velocities. This property is called *mixing*, and it is essential for the validity of the ergodic hypothesis.

We conclude that macroscopic properties are averages over the microscopic variables and give predictable values if the system is sufficiently large. One goal of statistical mechanics is to determine these averages and give a microscopic basis for the laws of thermodynamics. It is remarkable that these laws depend on the fact that the microscopic dynamics of the constituents of the macroscopic systems we encounter in our everyday experience are chaotic.

1.9 Phase Changes and Cooperative Effects

Most of our simulations of Lennard-Jones systems have so far been for dilute gases. What do you think would happen if we made the density higher? Would a system of interacting particles form a liquid or a solid if the temperature or the density were chosen appropriately? The existence of different phases is explored in Problem 1.9.

Problem 1.9. Different phases

(a) Program `LennardJonesMD` simulates an isolated system of N particles interacting via the Lennard-Jones potential. Choose $N=144$ and $L=18$ so that the density $\rho = N/L^2 \approx 0.44$. The initial positions are chosen at random except that no two particles are allowed to be closer than the length σ. Run the simulation and satisfy yourself that this choice of density and resultant total energy corresponds to a gas. What are your criteria?

(b) Slowly lower the total energy of the system. (The total energy is lowered by rescaling the velocities of the particles.) If you are patient, you will be able to observe "liquidlike" regions. How are they different from "gaslike" regions?

(c) If you decrease the total energy further, you will observe the system in a state roughly corresponding to a solid. What are your criteria for a solid? Explain why the solid that we obtain in this way will not be a perfect crystalline solid.

(d) Describe the motion of individual particles in the gas, liquid, and solid phases.

(e) Speculate on why a system of particles interacting via the Lennard-Jones potential can exist in different phases. Is it necessary for the potential to have an attractive part for the system to have a liquid phase? Is the attractive part necessary for there to be a solid phase? Describe a simulation that would help you answer this question. □

It is remarkable that a system with the same interparticle interaction can be in different phases. At the microscopic level, the dynamics of the particles is governed by the same equations of motion. What changes? How does a phase change occur at the microscopic level? Why doesn't a liquid crystallize immediately after its temperature is lowered quickly? What happens when it does begin to crystallize? We will find in later chapters that a phase change is an example of a *cooperative* effect. Familiar examples of phase transitions are the freezing and boiling of water. Another example with which you might be familiar is the loss of magnetism of nickel or iron above a certain temperature (358°C for nickel). Other examples of cooperative effects are the occurrence of gridlock on a highway when the density of vehicles exceeds a certain value, and the occurrence of an epidemic as a function of immune response and population density.

1.10 Models of Matter

There are many models of interest in statistical mechanics, corresponding to the wide range of macroscopic systems found in nature and made in the laboratory. So far we have discussed a simple model of a classical gas and used the same model to simulate a classical liquid and a solid.

One approach to understanding nature is to develop models that can be understood theoretically but are rich enough to show the same qualitative features that are observed in nature. We will emphasize models of this type in our discussions. In some contexts these models can yield quantitative data that approximates data from experiments. The models that we will consider include the following.

1.10.1 The ideal gas

The simplest models of macroscopic systems are those for which there are no interactions between the individual constituents of the system. For example, if a system of particles is very dilute, collisions between the particles will be rare and can be neglected under most circumstances. In the limit that the interactions between the particles can be neglected completely, the system can be modeled as an *ideal gas*. The ideal classical gas allows us to understand much about the behavior of dilute gases, such as those in the Earth's atmosphere. The quantum versions will be useful in understanding blackbody radiation, electrons in metals, and the low temperature behavior of crystalline solids.

The historical reason for the term "ideal" is that the neglect of interparticle interactions allows us to do calculations analytically. However, the neglect of interparticle interactions raises other issues. For example, how does an ideal gas reach equilibrium if there are no collisions between the particles?

1.10.2 Interparticle potentials

A common form of the potential between two neutral atoms is the Lennard-Jones potential given in (1.1) and Figure 1.1. This potential is the usual choice for studies where the focus is on fundamental issues, rather than on the properties of a specific material.

An even simpler interaction is purely repulsive and is given by

$$u(r) = \begin{cases} \infty & (r \le \sigma) \\ 0 & (r > \sigma). \end{cases} \tag{1.3}$$

A system of particles interacting via (1.3) is called a system of hard spheres, hard disks, or hard rods, depending on whether the spatial dimension is three, two, or one, respectively. The properties of dense gases and liquids will be discussed in Chapter 8.

1.10.3 Lattice models

For another class of models, the particle positions are restricted to a lattice or grid and the momenta of the particles are irrelevant. In the most popular model of this type, the "particles" correspond to magnetic moments. At high temperatures the magnetic moments are affected by external magnetic fields, but the interaction between the moments can be neglected.

The simplest nontrivial lattice model that includes interactions is the *Ising model*, the most important model in statistical mechanics. The model consists of magnetic moments or spins which are located on the sites of a lattice such that each spin can take on one of two values designated as up and down or ± 1. The interaction energy between two neighboring spins is $-J$ if the two spins point in the same direction and $+J$ if they point in opposite directions. One reason for the importance of this model is that it is one of the simplest to have a phase transition, in this case, a phase transition between a ferromagnetic state and a paramagnetic state. The Ising model is discussed in Chapter 5.

We will focus on three classes of models—the ideal classical and quantum gas, classical systems of interacting particles, and the Ising model. These models will be used in many contexts to illustrate the ideas and techniques of statistical mechanics.

1.11 Importance of Simulations

Only simple models (such as the ideal gas) or special cases (such as the one- and two-dimensional Ising models) can be solved exactly by analytical methods. Physicists frequently use simple models. For example, ignoring air friction allows us to learn much about the trajectory of a baseball in flight. However, if we want to understand the trajectory of a curve ball, we need a model that includes the interactions of a baseball with the air. The same is true in statistical mechanics. We can learn much from simple models, but we need more complex models to understand certain properties of materials. Usually, analytical calculations are insufficient to find the consequences of more complex models, and computational tools are needed.

Statistical mechanics has grown in importance over the past several decades because computers and new algorithms have made it possible to explore the behavior of more complex systems. As our models become more realistic, it is likely that they will require the computer for understanding many of their properties. Frequently the goal of a simulation is to explore the qualitative behavior of a model, so that we have a better idea of what type of theoretical analysis might be possible and what type of laboratory experiments should be done. Simulations also allow us to compute many quantities, some of which cannot be measured in a laboratory experiment. Theory, experiment, and simulations each play an important and complementary role in understanding nature.

Not only can we simulate reasonably realistic models, we also can study models that are impossible to realize in the laboratory, but are useful for providing a deeper theoretical understanding of real systems. For example, simulations of the Ising model and systems of particles with the interaction in (1.3) in four, five, and six dimensions have helped us understand the behavior of the three-dimensional system.

Simulations cannot replace laboratory experiments and are limited by the finite size of the systems and by the short duration of our runs. For example, at present, the longest simulations of simple liquids are for no more than 10^{-6} s.

1.12 Dimensionless Quantities

The different units used in science can be confusing. One reason is that sometimes the original measurements were done before a good theoretical understanding was achieved. For example, the calorie was created as a unit before it was understood that heat transfer was a form of energy transfer. Even today we frequently become confused using small calories and big calories and converting each to joules.

It is frequently convenient to use dimensionless quantities. These quantities can be defined by taking the ratio of two quantities with the same units. For example, the measure of the angle θ in radians is the ratio of the arc length s on a circle subtended by the angle to the radius r of the circle: $\theta = s/r$.

It is also useful to have another quantity with the same dimensions to set the scale. For example, if the speed v of a particle is close to the speed of light c, we usually specify the speed as the ratio v/c. (You might recognize the notation $\beta \equiv v/c$.) The use of dimensionless variables makes the relevant equations simpler and make it easier to perform algebraic manipulations, thus reducing the chance of errors. Another reason to use dimensionless variables is to make it easier to evaluate the importance of quantities in making approximations. For example, if $\beta \ll 1$, we know that relativistic corrections are not needed in most contexts.

We will frequently consider the high and low temperature behavior of a thermal system. What characterizes high temperature? To answer this question, we need to find a typical energy ϵ in the system and consider the dimensionless ratio $\tilde{T} \equiv kT/\epsilon$. An example is the behavior of electrons in metals at room temperature. We will find in Chapter 6 that it is convenient to choose ϵ to be the maximum kinetic energy of an electron at zero (absolute) temperature. In quantum mechanics, no two electrons can be in the same state, and hence ϵ is nonzero. In this case, if $\tilde{T} \ll 1$, the temperature of the system is low (even though the system is at room temperature).

Although $\tilde{T}$ is dimensionless, it is frequently convenient to call $\tilde{T}$ a temperature and sometimes even to denote it by T, with the understanding that T is measured in terms of ϵ. We have already done so in the context of Problems 1.5 and 1.7.

Another reason to use dimensionless quantities is that computers do not easily manipulate very small or very large numbers. Thus, it is best that all quantities be within a few orders of magnitude of 1. In addition, the use of dimensionless quantities allows us to do a simulation or an analytical calculation that is valid for many systems. To find the value of a quantity for a specific system, we just multiply the dimensionless quantity by the relevant quantity that sets the scale for the system of interest.

1.13 Summary

In principle, we have discussed most of the important concepts in thermodynamics and statistical physics, but it will take you a while before you understand these concepts in depth. Your understanding of these concepts and the methods of statistical and thermal physics will increase as you work with these concepts in different contexts. However, there is no unifying equation such as Newton's second law of motion in mechanics, Maxwell's equations in electrodynamics, and Schrödinger's equation in nonrelativistic quantum mechanics. The concepts are universal, but their application to particular systems is not.

Thermodynamics might not seem to be as interesting to you when you first encounter it. However, an understanding of thermodynamics is important in many contexts, including societal issues such as the Industrial Revolution, global climate change, and the development of alternative energy sources.

Thermodynamics and statistical mechanics have historically been applied to gases, liquids, and solids. This application has been very fruitful and is one reason that condensed matter physics, materials science, and chemical physics are rapidly evolving and growing areas. Examples of new systems of interest include high temperature superconductors, low dimensional magnets and conductors, composites, and biomaterials. Scientists are also taking a new look at more traditional condensed systems, such as water and other liquids, liquid crystals, and polymers. In addition to interest in macroscopic systems, there is growing interest is *mesoscopic* systems, systems that are neither microscopic nor macroscopic, but are in between, that is, between $\sim 10^2$ and $\sim 10^6$ particles.

The techniques and ideas of statistical mechanics are now being used outside traditional condensed matter and chemical physics. The field theories of high energy physics, especially lattice gauge theories, use the methods of statistical mechanics. New methods of doing quantum mechanics convert calculations to path integrals that are computed numerically using methods of statistical mechanics. Theories of the early universe use ideas from thermal physics. For example, we speak about the universe being quenched to a certain state in analogy to materials being quenched from high to low temperatures. We already have seen that chaos provides an underpinning for the need for probability in statistical mechanics. Conversely, many of the techniques used in describing the properties of dynamical systems have been borrowed from the theory of phase transitions, an important area of statistical mechanics.

In recent years statistical mechanics has evolved into the more general field of *statistical physics*. Examples of systems of interest in the latter area include earthquake faults, granular matter (for example, sand), neural networks, active matter (for example, flocks of birds), porous media such as rocks, models of computing, and the distribution of wealth in a society. Statistical physics is characterized more by its techniques than by the problems that are of interest. This universal applicability makes the techniques more difficult to understand, but makes the journey more exciting.

1.14 Supplementary Notes

1.14.1 Approach to equilibrium

In Problem 1.2 we found that $n(t)$, the number of particles on the left side of the box at time t, decreases in time from its initial value to its equilibrium value in an almost deterministic manner if $N \gg 1$. It is instructive to derive the time dependence of $n(t)$ to show explicitly how chance can generate deterministic behavior.

We learned that, if we run the simulation once, $n(t)$ will exhibit fluctuations and not decay monotonically to equilibrium. Suppose that we do the simulation many times and average the results of each run at a given time t. As discussed in Section 1.8, this average is an ensemble average, which we will denote by $\overline{n}(t)$. If there are $\overline{n}(t)$ particles on the left side after t moves, the change in $\overline{n}$ in the time interval Δt is given by

$$\Delta\overline{n} = \left[\frac{-\overline{n}(t)}{N} + \frac{N - \overline{n}(t)}{N}\right]\Delta t, \tag{1.4}$$

where Δt is the time between moves of a single particle from one side to the other. Equation (1.4) is equivalent to assuming that the change in $\overline{n}$ in one time step is equal to the probability that a particle is removed from the left plus the probability that it is added to the left. (In the simulation we defined the time so that the time interval Δt

between changes in $n(t)$ was equal to 1.) If we treat $\overline{n}$ and t as continuous variables and take the limit $\Delta t \to 0$, we have

$$\frac{\Delta \overline{n}}{\Delta t} \to \frac{d\overline{n}}{dt} = 1 - \frac{2\overline{n}(t)}{N}. \tag{1.5}$$

The solution of the differential equation (1.5) is

$$\overline{n}(t) = \frac{N}{2}\left[1 + e^{-2t/N}\right], \tag{1.6}$$

where we have used the initial condition $\overline{n}(t=0) = N$. We see that $\overline{n}(t)$ decays exponentially to its equilibrium value $N/2$. How does this form (1.6) compare to your simulation results for various values of N?

From (1.6) we can define a *relaxation time* τ as the time it takes the difference $[\overline{n}(t) - N/2]$ to decrease to $1/e$ of its initial value. Because $\tau = N/2$, $\overline{n}(t)$ for large N varies slowly, and we are justified in rewriting the difference equation (1.4) as a differential equation.

Problem 1.10. Independence of initial conditions
Show that if the number of particles on the left-hand side of the box at $t = 0$ is equal to $n(0)$ rather than N, the solution of (1.5) is

$$\overline{n}(t) = \frac{N}{2} - \frac{N}{2}\left[1 - \frac{2n(0)}{N}\right] e^{-2t/N}. \tag{1.7}$$

Note that $\overline{n}(t) \to N/2$ as $t \to \infty$ independent of the value of $n(0)$. □

1.14.2 Mathematics refresher

As discussed in Section 1.13, there is no unifying equation in statistical mechanics to be solved in a variety of contexts. For this reason we will use many mathematical tools. Section 2.24.1 summarizes the mathematics of thermodynamics, which makes much use of partial derivatives. The Appendix summarizes some of the mathematical formulas and relations that we will use.

If you can do the following problems, you have a good background for much of the mathematics that we will use in the following chapters.

Problem 1.11. Common derivatives
Calculate the derivative with respect to x of the following functions: e^x, e^{3x}, e^{ax}, $\ln x$, $\ln x^2$, $\ln 3x$, $\ln 1/x$, $\sin x$, $\cos x$, $\sin 3x$, and $\cos 2x$. □

Problem 1.12. Common integrals
Calculate the following integrals:

$$\int_1^2 \frac{dx}{2x^2}, \qquad \int_1^2 \frac{dx}{4x}, \qquad \int_1^2 e^{3x}\, dx, \qquad \int x^{-\gamma}\, dx. \tag{1.8}$$

□

Problem 1.13. Partial derivatives
Calculate the partial derivatives of $x^2 + xy + 3y^2$ with respect to x and y. ☐

Problem 1.14. Taylor series approximations
Calculate the first three nonzero terms of the Taylor series approximations about $x = 0$ for the following functions:

$$e^{ax}, \qquad \ln(1+x), \qquad (1+x)^n, \tag{1.9}$$

where a and n are constants. ☐

Vocabulary

thermodynamics, statistical mechanics
macroscopic system, microstate, macrostate
specially prepared state, most probable macrostate
equilibrium, fluctuations, relative fluctuations
thermal contact, temperature
sensitivity to initial conditions, chaos
models, computer simulations, molecular dynamics
dimensionless variables

Additional Problems

Problem 1.15. The dye is cast

(a) What would you observe if a small amount of black dye is placed in a glass of water?
(b) Suppose that a video were taken of this process and the video run backward without your knowledge. Would you be able to determine whether the video was being run forward or backward?
(c) Suppose that you could watch a video of the motion of an individual ink molecule. Would you be able to determine if the video was being shown forward or backward? ☐

Problem 1.16. Fluid as metaphor
Why is "heat" treated as a fluid in everyday speech? After all, most people are not familiar with the caloric theory of heat. ☐

Problem 1.17. Do molecules really move?
Cite evidence from your everyday experience that the molecules in a glass of water or in the surrounding air are in seemingly endless random motion. ☐

Problem 1.18. Temperature
How do you know that two objects are at the same temperature? How do you know that two bodies are at different temperatures? ☐

Problem 1.19. Time-reversal invariance
Show that Newton's equations are time-reversal invariant. ☐

Problem 1.20. Properties of macroscopic systems
Summarize your present understanding of the properties of macroscopic systems. ☐

Problem 1.21. What's in a name?
Ask some of your friends why a ball falls when released above the Earth's surface. Then ask them what makes rolling balls come to rest. Are the answers of "gravity" and "friction" satisfactory explanations? What would be a more fundamental explanation for each phenomena? ☐

Problem 1.22. Randomness
What is your present understanding of the concept of "randomness"? Does "random motion" imply that the motion occurs according to unknown rules? ☐

Problem 1.23. Meaning of abstract concepts
Write a paragraph on the meanings of the abstract concepts "energy" and "justice." [See Feynman, Leighton, and Sands, Vol. 1, Chap. 4, for a discussion of why it is difficult to define such abstract concepts.] ☐

Problem 1.24. Bicycle pump
Suppose that the handle of a plastic bicycle pump is rapidly pushed inward. Predict what happens to the temperature of the air inside the pump and explain your reasoning. ☐

Problem 1.25. Granular matter
A box of glass beads is an example of a macroscopic system if the number of beads is sufficiently large. In what ways is such a system different from the macroscopic systems such as a glass of water that we have discussed in this chapter? ☐

Suggestions for Further Reading

P. W. Atkins, *The Second Law*, Scientific American Books (1984). A qualitative introduction to the second law of thermodynamics and its implications.

Manfred Eigen and Ruthild Winkler, *How the Principles of Nature Govern Chance*, Princeton University Press (1993).

Richard Feynman, R. B. Leighton, and M. Sands, *Feynman Lectures on Physics*, Addison-Wesley (1964). Volume 1 discusses the nature of energy and work.

Martin Goldstein and Inge F. Goldstein, *The Refrigerator and the Universe*, Harvard University Press (1993).

Harvey Gould, Jan Tobochnik, and Wolfgang Christian, *An Introduction to Computer Simulation Methods*, 3rd ed., Addison-Wesley (2006). See Chap. 8 for a discussion of molecular dynamics. This text can be freely downloaded from <compadre.org> and purchased from Amazon.

Robert H. Romer, "Heat is not a noun," Am. J. Phys. **69**(2), 107–109 (2001). See also Art Hobson, "The language of physics," Am. J. Phys. **69**(6), 634 (2001); David T. Brookes and Eugenia Etkina, "Using conceptual metaphor and functional grammar to explore how language used in physics affects student learning," Phys. Rev. ST Phys. Educ. Res. **3**(1), 010105-1–16 (2007).

A listing of many of the textbooks on statistical mechanics and thermodynamics can be found at <www.compadre.org/stpbook/>. Some of our favorites are listed in the following. It is

a good idea to look at several books while you are learning a subject. Sometimes the same argument with slightly different wording can seem clearer.

Peter Atkins, *The Laws of Thermodynamics: A Very Short Introduction*, Oxford University Press (2010).

Stephen Blundell and Katherine Blundell, *Thermal Physics*, 2nd ed., Oxford University Press (2006).

Craig F. Bohren and Bruce A. Albrecht, *Atmospheric Thermodynamics*, Oxford University Press (1998).

F. Mandl, *Statistical Physics*, 2nd ed., John Wiley & Sons (1988).

F. Reif, *Statistical Physics*, Vol. 5 of the Berkeley Physics Series, McGraw-Hill (1967). This text was probably the first to make use of computer simulations to explain some of the basic properties of macroscopic systems.

F. Reif, *Fundamentals of Statistical and Thermal Physics*, McGraw-Hill (1965). This classic text is now published by Waveland Press. Our presentation owes a large debt to the two books by Reif on statistical physics.

Daniel V. Schroeder, *An Introduction to Thermal Physics*, Addison-Wesley (2000).

Index