Many-particle Systems, 3

Bare essentials of statistical mechanics

Atoms are examples of many-particle systems, but atoms are extraordinarily simpler than *macroscopic* systems consisting of 10^{20} - 10^{30} atoms. Despite their great size, many properties of macroscopic systems depend intimately on the microscopic behavior of their microscopic constituents. The proper quantum mechanical description of an *N*-particle system is a wavefunction that depends on 3N coordinates (3 ways of moving, in general, for every particle) and 4N quantum numbers (3 motional quantum numbers and 1 spin quantum number for every particle). (If the "particles" are molecules there might be additional quantum numbers describing rotations and internal vibrations.) The latter is referred to as a **microstate**. In a macroscopic solid or liquid, there are about 10^{23} atoms per cm³; in a gas, there are about 10^{20} . Thus, the amount of information required to specify the wavefunction for such a system is scandalously large. A fundamental theory of macroscopic systems cannot be based on all of the microscopic details. Instead, a new way of thinking is required. This is known as **statistical mechanics**.

Important to the success of statistical mechanics is the recognition that the system interacts with its surroundings. These interactions are typically so frequent that it is impossible to keep track of them, so a simplifying assumption is invoked: the interactions cause *random* (see the Appendix below) transitions between the microstates. Thus, if a system is initially prepared in a unique microstate, in short order (about 10⁻¹³ s or less), the system will be in another and we can't tell which it will be. This kind of uncertainty is different from quantum uncertainty: it's due to our overwhelming ignorance of all of the possible information.

Example: Consider a standard deck of 52 playing cards. A microstate is the actual sequence of individual cards from the top (card #1) to the bottom (card #52). Transitions are random rearrangements: "shufflings." Card #1 might be any one of the 52 cards; card #2 might be any one of the remaining 51; ... card #52 can only be the one card left over after all the others have found a home. That means there are $52! = 8 \times 10^{67}$ different microstates. Suppose the deck starts out in some highly ordered state (according to a human observer who happens to value that arrangement)—such as king of spades, gueen of spades, ..., ace of spades, king of hearts, ... ace of diamonds, perhaps the order of a brand-new deck. Let's define a "shuffle" as follows: a small number of cards in the deck are chosen at random, removed from the deck, then replaced at randomly chosen positions. If a shuffle doesn't mix too much, one shuffle will produce a new microstate with less order, though guite a bit will remain. A second shuffle typically produces less order still. The more shuffles the less order, on average. Of course, any sequence of random shuffling *might* result in reproducing the initial microstate, but on average that should happen only about 1 out of 8x10⁶⁷ times. That's a lot less likely than winning the Powerball Lottery. The number of possible microstates, 8x10⁶⁷, is a staggeringly large number. If every shuffle produced a different microstate it would take at least 8x10⁶⁷ shuffles to explore all possible microstates. It would take a billion high-speed computers shuffling at the rate of 10¹³ per second (the rate of molecular collisions in a glass of water) 8x10⁵⁴ seconds to complete the task (assuming no two computers make the same shuffling order)—that's about 2.5x10⁴⁷ years!!! (The age of the universe is estimated to be only 1.4x10¹⁰ years; clearly, during the whole of human card playing history only a very small fraction of all of the possible deck orders has yet been produced.)

The goal of statistical mechanics is to describe how macroscopic variables are related to averages over many microscopic states. Thus, for example, we might like to know how the voltage of a thermocouple in a glass of water is related to the quantum states of the water molecules. A temperature reading in the laboratory might take seconds, but in that time 10¹³ different microstates might occur. The thermocouple reading therefore results from an average over a huge number of quantum states. If the thermocouple reading and other macroscopic measurements are approximately constant, changing by only small fluctuations from second-to-second, the system is said to be in **thermal equilibrium**. At the micro-level that means that the sequence of microstates contributing to the macroscopic measurement is "typical" of all of the possible microstate sequences. That, in turn, occurs when the probability of the system achieving each of its microstates does not change in time, a condition called **statistical equilibrium**. When the system is in statistical equilibrium an average over all possible microstates (a calculation) produces roughly the same result as an average over a *representative* sample of microstates (a measurement).

Example: Suppose you are playing poker and wonder what the chances are of drawing a flush—any five cards of the same suit. Say you are dealt the first five cards from the deck. If the deck is in the initially ordered state described in the previous example, you're in great shape: you get a king-high straight flush! After one shuffle, you might still see quite a few common suit cards; after another shuffle, possibly fewer, and so on. It's reasonable to expect your chances of a flush to drop from certainty to a steady value after a number of shuffles (as statistical equilibrium is established). So, what is that steady value? A flush poker hand is a kind of macrostate: it doesn't require specification of the complete microstate of the deck. If the deck has been shuffled into a condition of statistical equilibrium, then the probability of each of its 8x10⁶⁷ microstates is equal. We can employ statistical equilibrium to calculate the probability of a flush. The first card you draw determines what suit the flush has to be in and depletes the pool of suit cards by 1 and the total pool by 1 also. Thus, the probability the second card is of the same suit is 12/51. Similarly, the probabilities for the third, fourth, and fifth cards being in the same suit as the first card are 11/50, 10/49, and 9/48, respectively. Note that these probabilities decrease. Future events (cards) are correlated with past events (cards). The total probability is the product of the individual values, about 0.002. Thus, if you play billions of times you should expect to get a flush hand about once every 500 hands.

Statistical mechanics often invokes another approximation, namely that the system's particles *do not interact with one another except for boson attraction or fermion repulsion.* That is, they don't interact via electromagnetic or other "classical" forces. Such a system is referred to as an *ideal gas*. The most profound consequences of quantum mechanics in the macroscopic world arise from the indistinguishability of identical particles. We focus hereafter on that situation; all of the systems we examine consist of particles of one kind that interact only through the symmetry or antisymmetry of their wavefunctions. The system wavefunction can then be separated into a symmetric (for bosons) or antisymmetric (for fermions) combination of products of *N* single-particle wavefunctions, each depending on 3 coordinates and 4 quantum numbers (such as, n_x, n_y, n_z, m_z for a particle with spin component

 m_s in a 3D infinite well–a *single-particle state*), for structureless particles—which we also restrict to in the following. A system of *N* identical, *noninteracting* particles has total energy

$$E_{S} = \sum_{all \sigma} N_{\sigma} \varepsilon_{\sigma}$$

where the *S* represents a system microstate (that is, 4N quantum numbers) and the "little σ " means a single-particle state consisting of 4 quantum numbers; ε_{σ} is the energy of the single-

particle state and N_{σ} is the number of particles in that state. The total number of particles in the system is

$$N = \sum_{all \ \sigma} N_{\sigma} \, .$$

Generally, both *E* and *N* might change due to interactions of the many-particle system with its surroundings (even though there are no internal interactions). Statistical mechanics is interested in the average values $\overline{E} = \overline{E}_s = \sum_{\sigma} \overline{N}_{\sigma} \varepsilon_{\sigma}$ and $\overline{N} = \sum_{\sigma} \overline{N}_{\sigma}$, which are calculated by invoking statistical equilibrium. In fact, all of the interesting statistical mechanical properties of systems of indistinguishable, noninteracting particles are calculated by knowing the appropriate \overline{N}_{σ} .

Note that a small change in \overline{E} is expressed as

$$d\overline{E} = \sum_{\sigma} (\varepsilon_{\sigma} d\overline{N}_{\sigma} + \overline{N}_{\sigma} d\varepsilon_{\sigma});$$

moreover,

$$d\overline{N} = \sum_{\sigma} d\overline{N}_{\sigma}$$

The latter can be divided into two parts: $\sum_{\sigma} [(d\bar{N}_{\sigma})_{\chi} + (d\bar{N}_{\sigma})_{T}]$, where, by definition, $\sum_{\sigma} (d\bar{N}_{\sigma})_{\chi} = 0$.

The cartoon to the right distinguishes between the "X" and "T" processes. The lines represent energy levels for states of single-particles and the dots indicate which level is occupied. In the

top row the sum of the numbers of particles in all states is three before and after the change. The energy difference between the two microstates is due only to energy exchange between the system and its surroundings, and is associated with *eXcitations* (or de-excitations) (X). In the bottom row a new particle is added (without the excitation levels of the original particles changing) and the sum of the total number

original particles changing) and the sum of the total number $\xrightarrow{}$ $\xrightarrow{}$

$$d\overline{E} = \sum_{\sigma} [(d\overline{N}_{\sigma})_{X} \varepsilon_{\sigma} + (d\overline{N}_{\sigma})_{T} \varepsilon_{\sigma} + \overline{N}_{\sigma} d\varepsilon_{\sigma}].$$

While $\sum_{\sigma} (d\overline{N}_{\sigma})_{\chi}$ is zero, $\sum_{\sigma} (d\overline{N}_{\sigma})_{\chi} \varepsilon_{\sigma}$ need not be zero, as in the cartoon above, for example. Assuming that the system has neither net center-of-mass motion nor coherent motion about the

Assuming that the system has neither net center-of-mass motion nor coherent motion about the center-of-mass (such as a collective rotation or vibration), \overline{E} is the system's *internal energy*. The right hand side indicates that internal energy can be changed in three ways: (1) the term $\sum_{\sigma} (d\overline{N}_{\sigma})_{\chi} \varepsilon_{\sigma}$ represents energy change by "exchange of excitations" only; such a process is

called *heat* and is usually written dQ (heat is energy exchanged due to temperature differences only); (2) the term $\sum_{\sigma} (d\bar{N}_{\sigma})_{T} \varepsilon_{\sigma}$ represents energy change by "exchange of particle

number" only; such a process is called *chemical work* and is usually written $\mu d\overline{N}$, where μ is the system's *chemical potential* (note: **particles diffuse at constant temperature—if they're not blocked from doing so—from higher chemical potential to lower**; chemical potential is a measure of how reluctant a system is to accepting particles from its surroundings [the previous example suggests that, in general, bosons have lower chemical potentials than



fermions]); (3) the term $\sum_{\sigma} \bar{N}_{\sigma} d\varepsilon_{\sigma}$ represents energy change by "changing level spacing"

only; this process involves changing the system's volume or manipulating external (electric and magnetic) fields and is called (physical) *work* and is written -dW (where the minus sign is because work done *by* the system is defined as > 0 and it reduces the internal energy). In summary,

$$d\overline{E} = dQ + \mu d\overline{N} - dW$$

which is the classical *First Law of Classical Thermodynamics*. In other words, classical thermodynamics is really just quantum mechanics plus random scrambling of states.

Appendix: Chaos and randomness

A *dynamical system* consists of a state and a rule for changing the state in time. In a deterministic dynamical system the future is completely determined by the past. The dimension of the dynamical system is related to how many states in the past are necessary to fix the next state in the future: a one-dimensional dynamical system only requires one past state to fix the future, a six-dimensional system requires six past states. In the 1960s, Edward Lorenz, a meteorology professor at MIT discovered by accident that dynamical systems with nonlinear update rules could produce sequences of states that looked unpredictable; this behavior is called (deterministic) chaos. Lorenz probably ensured that many others would study chaos in subsequent years because of his phrase, "The Butterfly Effect." The Butterfly Effect occurs when a chaotic system is started in two different, but very close, states; at first, the updated states stay close, but after awhile they depart from each other. Lesson: if you want to predict the output of a chaotic system and have uncertain initial data you might be able to predict a short time into the future, but after awhile your ability to predict will be almost totally lost. In fact, the higher the dimension of the chaotic dynamical system the faster prediction will fail, if the initial data are a little uncertain. If the initial data are perfectly clean and precise, then it takes at least as many initial states as the dimension of the system to predict what happens in the future. A *random* process, in principle, produces a sequence of states the next value of which is unpredictable no matter how much initial data one has. In this sense, a random process is an infinite dimensional chaotic system. Of course, if the system of interest is just high dimensional enough, its output is probably as good as random. The "random number generator" in your computer, for example, is not random at all; it's a high dimensional chaotic system. In fact, because testing if a process is random requires an infinite number of states we're not sure whether anything is precisely random (though we speculate that quantum behavior might be).