Physics 3700

Introduction. Equilibrium. Temperature. Heat, Work, and The First Law.

Relevant sections in text: §1.1–1.5

The Big Picture

"Thermal Physics" is about understanding the behavior of physical systems focusing on the ideas of heat and temperature. A lot of this course will be devoted to understanding heat and temperature at various levels of sophistication. The scope of thermal physics is remarkable: it pretty much applies to all "macroscopic" systems – and we shall often use thermal physics insights to better understand microscopic physics! One can use thermal physics to understand something as simple as a melting ice cube or something as exotic as a black hole. We shall see that the "thermodynamic" behavior of a macroscopic system is ultimately due to the fact that it consists of many, many subsystems.

Thermal physics can quickly become confusing unless we make some concepts clear from the outset, namely, *observables*, *states*, and *laws*. These concepts are in fact the fundamental constituents of any theory. A theory describing/explaining a given physical system will include a number of ingredients, but almost always it includes a notion of observables, states and laws. It will be very helpful for your sanity in this course to understand clearly what these terms mean, so let's spend a moment on this.

Observables

Observables are, of course, things you can measure about the system (or infer from measurements). Understanding and predicting the values of the observables which occur in this or that situation is the principal goal of the physicist. The precise mathematical representation of observables will vary from theory to theory. In thermodynamics the observables are represented by real variables – more precisely, real-valued functions of the state, as will be explained. (Contrast this with quantum mechanics where observables are represented by linear operators on a vector space.) Examples of observables in thermodynamics include temperature, pressure, volume, entropy, *etc.* Because thermodynamics concerns itself principally with macroscopic quantities (*e.g.*, energy of an atom). It's not that you can't measure the energy of an atom, or that we won't consider such things, it's just not typically what you are focusing on in the thermodynamical setting. Unless we explicitly state otherwise, we shall reserve the term "observable" only for the macroscopic quantities.

States

The idea of a state is a little more abstract. Roughly speaking, the state characterizes what the system is doing at a given time. How do you know what the system is doing? By measuring its observables! Knowing the state of the system at a given time is equivalent to knowing the outcome of all possible measurements at that time.[†] So, for example, given the state of a thermodynamic system one is able to say what is the volume, pressure, entropy, and so on. Normally, measuring a sufficient number of observables (not necessarily all of them) will tell you what is the state of the system, that is, it will allow you to infer the values of all other observables via your model for the system (as far as possible^{*}). Sometimes the minimum number of observables needed to determine a state is called the "dimensionality of the phase space" - or just the "dimensionality" - or even the "dimension" - of the system. One of the main reasons thermodynamics is successful as a way to describe nature is that the dimensionality of typical macroscopic systems is not too high (despite the huge number of microscopic constituents). Of course, states of a system will change in time for various reasons. Different states of a system are distinguished by a difference in at least one of the things you can measure. If any one observable for a system changes its value, then the state of the system has changed. Conversely, if the state of a system changes, at least one of its observables changes value.

Laws

You are familiar with the idea of "laws" from Newtonian mechanics. Roughly speaking, laws specify how we relate the observables to the states, how the states vary as the observables change, and how the states vary in response to external influences. In thermodynamics there are 4 laws (or 3, depending upon how you count them), as you may already know. These four laws have been summarized as: "You have to play the game. You can't win the game. You can't even break even. You cannot quit the game." It sounds a bit depressing, I suppose. Unfortunately, thermodynamics is a very accurate physical description of the universe. Anyway, our principal task in this course is to understand what these laws really mean and how to use them to understand (specific models of) various physical systems.

The four laws of thermodynamics are meant to apply to any (macroscopic) system. There are also rules – sometimes called laws, more often called *equations of state* – which

[†] Note that this use of the term "state" has nothing to do with the usage which refers to different "phases" of matter (solid, liquid gas,...). We shall try to consistently use "state" to refer the outcome of measurements meaning and "phase" to distinguish things like solid, liquid, gas.

^{*} I keep adding this parenthetical remark because there are theories like quantum mechanics where one is often only able to obtain statistical information about the values of observables.

pertain to (models of) specific physical systems. An equation of state gives a relation between the observables which arises because of the particular details of the physical system being considered.* In thermodynamics, equations of state are empirically obtained or are derived from microscopic theories (see below). As an important example, you will recall the ideal gas law from elementary physics and/or chemistry: PV = NkT. Here P is the pressure in a gas, V is the volume of the gas, N is the number of molecules making up the gas, T is the absolute temperature, and k is Boltzmann's constant (see your text for more on ideal gases). If we suppose the number of molecules is fixed, the ideal gas equation of state tells us that of the three observables (P, V, T) any two determine the third for this system. This reflects the fact that the dimensionality of a fixed quantity of an ideal gas is two: we can completely characterize the state of the ideal gas, say, by measuring P and Vor by measuring T and V, etc. All other observables (internal energy, entropy, etc.) are then determined. Note that the ideal gas law is injected into thermodynamics from some other theory, e.q., the kinetic theory of gases. We will generally use "equations of state" for such relations between observables for a specific physical system and reserve the term "laws" for the 4 laws of thermodynamics.

Equilibrium

Aside from the above, the other principal part of the "rules" of thermodynamics is the notion of *equilibrium*. Strictly speaking, the states and observables of thermal physics only pertain to equilibrium situations.[†] Depending upon how a given system interacts with its environment, equilibrium can come in a variety of different flavors and is characterized via a variety of observables. Here are some important examples.

Mechanical equilibrium: Here all forces are balanced. A system typically arrives at mechanical equilibrium by moving, often by adjusting its volume – the volume changes until mechanical equilibrium is obtained.

Diffusive equilibrium: Here molecules comprising the system have no net motion into or out of the system. A system typically arrives at diffusive equilibrium by adjusting the number of this or that species of molecule in the system – the particle number changes (by exchange with the system's environment) until diffusive equilibrium occurs.

Mechanical and diffusive equilibrium should be familiar to you from Newtonian mechanics and chemistry, respectively. The new equilibrium idea in thermal physics is:

^{*} $\overline{\text{It is these equations of state which make finite the dimensionality of a thermodynamic system.}$

[†] The first law of thermodynamics postulates the idea of thermal equilibrium, as we shall see.

Thermal equilibrium: A system arrives at thermal equilibrium by exchanging energy with its environment due to a difference in temperature – see the discussion of the zeroth law below. When such an exchange of energy between two systems is possible we say that the systems are in "thermal contact". This energy exchange is the central issue in thermal physics.

The "dynamics" in thermodynamics refers to a change of state — from some initial state to some final state. Thermodynamic states are *always* equilibrium states, so a change of state necessarily means that the system is not in equilibrium while the state is changing. Consequently, we generally cannot use thermodynamic states/observables to describe what is happening when the state is changing; we can only describe the initial and final equilibria.*

"Theory" versus "Model"

The "rules" of thermodynamics are clear and unambiguous. The application of these rules to specific physical systems is where the art and challenge of the physicist comes into play. One has to settle on a *model* of the physical system of interest which is simple enough to allow for a tractable application of the rules of the theory, but which is sufficiently complex to accurately represent the physical features we are interested in. We shall see a variety of levels of complexity in our models of various elementary physical systems depending upon which behavior we are trying to capture.

While it is just a matter of nomenclature, it is important to keep clear in your mind the distinction between "theory" and "model". This can be confusing since sometimes texts will take (what I call) a "model", say, of an atom, and call it a "theory" of the atom, when the theory is (in my nomenclature) really the set of rules called quantum mechanics. As an example of the model versus theory distinction I am trying to make, consider the Sun-Earth system in Newtonian mechanics. Newton's laws (including his law of gravitation) constitute the theory. If you are just interested in the main features of the orbit of the Earth, you can take both bodies as "point particles", put the Sun at rest (since it is so massive) and use the usual $1/r^2$ force law to figure out how the Earth moves. This is one possible model which uses the theory. If you are interested in more details concerning the shape of the orbit, including their slowly varying properties in time, you should treat both the bodies as having finite size, take account of the (very small) motion of the Sun, take account of the tides, take account of other planets, and so forth. You use the same theory here (Newton's laws), you just make the model more detailed.

^{*} As we shall see, in some situations we *can* use thermodynamic states/observables to describe a changing state — at least in an approximate, idealized way using "quasi-static processes".

The scope of "thermal physics"

Thermodynamics is a self-contained, standalone theory of the behavior of macroscopic systems interacting among themselves and/or with an environment where one includes the phenomena of heat and temperature. The theory (states, observables, laws) operates independently from the microscopic details of the system. This makes thermodynamics very powerful because the underlying microscopic behavior is generally very complex. By the same token, thermodynamics cannot tell us about the microscopic details of the systems being considered, which we are also interested in. And the need for equilibrium also eliminates a number of interesting physical processes (*e.g.*, phase transitions) from direct, detailed examination (not that one cannot analyze certain aspects of non-equilibrium processes using thermodynamics, just not everything you might want to study). The study of the microscopic description provided by thermodynamics is often called *statistical thermodynamics* since it invariably involves large numbers of subsystems (atoms, molecules, spins, *etc.*) whose bulk statistical properties are studied in thermodynamics. We shall certainly make a lot of contact with statistical thermodynamics in this course.

After all these introductory remarks, we are ready to start the systematic development of the subject.

Thermal Equilibrium. Temperature. The Zeroth Law.

Suppose you take a cold glass of water into a warm room. The water interacts with its environment (the room, its air, *etc.*) for a while – principally by an exchange of energy. Eventually the water settles into an equilibrium state characterized by a particular temperature, volume, pressure, *etc.* for the water. During this "settling down", usually called *relaxation*, one does not properly speak of the state of the water* (although it can be done approximately as we shall see). The situation where the water's thermodynamic observables are again fixed is called *thermal equilibrium*. Roughly speaking, thermal equilibrium is characterized by no net exchange of energy of the system (water) with its environment (room). The time it takes to reach equilibrium is the *relaxation time*. Although thermodynamics is most definitely concerned with changes in physical systems due to interaction with other systems or with some "environment", strictly speaking thermodynamics does not deal with what is happening during relaxation. It describes the initial and final states only, *e.g.*, the change in temperature of the water. (Again, with certain approximations one can do a little better.) The way in which we change the state of the system – by letting it "warm up", by compressing it, or whatever – will often be referred to as this or that

^{*} As an illustration, during relaxation the temperature of the glass of water is, strictly speaking, not defined since, for example, the water near the outside of the glass will be "warmer" than the water near the center.

process. For example we are referring to processes in this sense when we say, "consider a process in which the water in a glass comes to equilibrium at room temperature...", or "consider a process in which the gas is compressed to half its initial volume at constant temperature...".

The notion of temperature is intimately tied to the notion of thermal equilibrium. Consider the glass of water interacting with the air. If we like, the water and air can be viewed together as constituting "the system", or we can treat each separately as "the system". Initially the water is "cooler" than the air. The water "warms up" and the air "cools down" § until equilibrium is established. Once equilibrium is again established, it is easy to see that the states of the air and water have changed. Indeed, some observables pertaining to each part of the system will have changed from their initial values, e.g., the volume of the water will have increased slightly, the pressure of the air will have decreased slightly. But, of course, the principal observable that has changed during the relaxation process is the temperature. The idea of temperature is that its value will be the *same* for both the air and water once equilibrium has been established. Thus we have the zeroth law of thermodynamics, which essentially says: Any two systems in thermal equilibrium have the same temperature. This is essentially a postulate allowing for the existence thermal equilibrium, with temperature being the name given to any observable which characterizes thermal equilibrium.[†] Two systems, each in equilibrium, but isolated from each other – more precisely, not in "thermal contact" – may have different temperatures. When the two systems are put in thermal contact, energy will flow from the system with the higher temperature to the system with the lower temperature until the temperatures are equal.

The description of temperature and thermal equilibrium just given might seem like an overly pedantic description of pretty obvious stuff. But a good physicist takes nothing for granted and always asks "why?". And it is a good idea to ask things like: Why does temperature exist? Why does energy flow from high temperature to low temperature? It is possible to answer such questions using a statistical description of the microscopic constituents of the macroscopic system, and we shall do this. But we will not go there yet. Instead, in thermodynamics, we simply postulate the ideas of temperature and thermal equilibrium and carry on.

The zeroth law postulates the existence of a new observable, temperature, which is the same for any systems in thermal equilibrium. The postulate does not say how this temperature is to be determined/measured. Using a specific physical model, we can use

[§] Of course, we often choose to ignore this usually small change in the air, much as we may choose to ignore the motion of the Sun in the Earth-Sun gravitating system.

[†] Much as Newton's first law can be viewed as a postulate allowing for the existence of inertial reference frames. We need such reference frames because it is in such frames that the second law "F = ma" is meant to hold. Likewise, we need to define/postulate thermal equilibrium, because thermodynamic observables are only defined in equilibrium states.

equations of state to determine a temperature. For example, the product of pressure and volume of a fixed quantity of an ideal gas allows us to determine the temperature, T = PV/Nk. Put differently, two ideal gases which are in "thermal contact" will exchange energy until they reach thermal equilibrium. The ideal gas law tells us that, while the pressures, volumes and number of particles of the two gases may differ, at equilibrium the product of the respective pressures and volumes divided by the respective total number of molecules will be the same. The idea of the zeroth law is that something analogous to PV/Nk will exist for any thermodynamic system.

You can think of the temperature of a system as being related to how much energy is in the system, which is a pretty common attitude toward temperature. This attitude is okay as far as it goes, but it is still pretty far from the correct meaning of temperature. For example, recalling that temperature is the same for two systems in thermal equilibrium, you might surmise that (at equilibrium) the energy in a glass of water need not be the same as the energy of the surrounding room. So temperature cannot be blatantly *identified* with energy, *per se*. As we develop our thermodynamic technology we will get a better and better explanation of what temperature is all about — understanding the precise meaning of temperature is one of the principal goals of this course.

To summarize: The zeroth law is the assertion that when two objects interact thermally, energy flows from one system to the other until (thermal) equilibrium is attained at which point there is an observable – temperature – which is the same for both systems.

Heat and Work. The First Law.

From the thermodynamic point of view, when a system interacts with its environment and changes its state the most fundamental consideration is that of the energy transfer involved. In particular, given the zeroth law, and the role of temperature in deciding how thermal energy flows between systems, this transfer of energy can be immediately divided into two types. *Heat* refers to a transfer of energy due to the difference in temperature of the system and its environment. Think of the glass of water (just discussed) as the system and the air in the room as the environment. The principal type of energy transfer in that example is heat. *Work* refers to any other transfer of energy, *e.g.*, due to changing some of the other observables by one means or another. For example, if you stir the water in the glass, you are adding energy by mechanical means and this is work, not heat.

Transfer of energy by mechanical means is familiar to you, *e.g.*, from Newtonian mechanics. *Transfer of energy by heat is the principal new ingredient in thermodynamics*. The thermodynamic concept "heat" characterizes a complex interaction at the microscopic level in terms of a new, macroscopic form of energy transfer. Eventually we will get a very nice microscopic explanation of heat (via statistical analysis), but in light of the zeroth law we can use the concept of heat as is without any microscopic interpretation.

Here is a very important conceptual point to keep in mind: Heat and work are **not** observables for a thermodynamic system. They represent influences of the environment causing a change of state. This will become more clear when we write down the first law of thermodynamics. It is not that the energy transfer to/from a system cannot be measured – it can. It is that the way we define the words "heat" and "work" in thermodynamics has them being distinguished from observables which characterize a given equilibrium state.

The first law of thermodynamics is really just a statement of conservation of energy. More precisely, the first law of thermodynamics is a statement about how the energy observable changes when the state of a system changes due to heat or work. Let U denote the *internal energy* of the system in question. U is called "internal energy" to distinguish it from some other pieces of the energy we will discuss later, and to keep clear that we are only studying the energy which is accessible via thermodynamic processes.^{*} For example, the internal energy of a monatomic ideal gas is simply the kinetic energy of all the atoms making up the gas. Let us consider a change of state of the system due to a process involving heat Q and work W – both representing energy entering the system. The observable Uwill change by an amount denoted ΔU . The first law of thermodynamics simply says

$$\Delta U = Q + W.$$

Note that if ΔU is positive, then the internal energy of the system increased during the change of state. Thus positive values for Q and/or W reflect energy transfer *into* the system. Negative values of Q and/or W indicate energy leaving the system. Some references use other conventions for the sign of Q and/or W, so beware of that when you go reading literature on thermodynamics.

Heat, Work, the First Law: Examples

Let us consider some simple, illustrative examples of heat, work and the first law. We will use an ideal gas as the system because it's just about the simplest system around, but also because many results obtained here easily generalize to other, more complex systems.

Consider a cylindrical container full of an ideal gas consisting of N molecules. The gas is our thermodynamical system. Observables include temperature T, energy U, pressure P, and volume V.[†] The ideal gas law tells us that

$$PV = NkT$$

^{*} For example, one does not normally contemplate processes which change the rest energy of the electrons in the atoms making up the system. One normally is not interested in the kinetic energy associated with translational motion of the system as a whole, *etc.*

 $[\]dagger$ N can be an observable, too. But we will hold it fixed in the present discussion, so it is just some number.

and an elementary result from kinetic theory (see your text) tells us that the energy is given by

$$U = \frac{f}{2}NkT,$$

where f is the number of "degrees of freedom" of the gas molecules.** f = 3 for a monatomic ideal gas. This result, like the ideal gas law, is not derivable from the laws of thermodynamics but rather from a microscopic model of the system. We take these equations of state as *defining* the ideal gas as a macroscopic system.

The energy U arises from the fact (to be proved later) that the average energy of a molecule in the gas is $\frac{f}{2}kT$ where f is the number of "degrees of freedom" of the molecule (to be defined later). Since there are so many particles in the gas, the thermal energy is essentially the number of particles times the average energy per particle. Note that this energy doesn't include things like, chemical bonding energy in the molecules, or the center of mass kinetic energy of the system as a whole, etc. For the processes of interest (for now, anyway) these are not changing and so we can adjust our zero point of energy to eliminate them. As usual (in non-gravitational physics, anyway), it is *changes* in energy (or lack thereof) which have significance.

We suppose one end of the cylindrical container can be moved in and out to adjust the volume. We have a "piston". We can do *compressional work* on the gas by moving the piston. This idea of compressional work is going to be pretty ubiquitous; it is the most common form of work we will consider. So we will now study compressional work here in some detail in the context of the ideal gas.

Let x denote the displacement of the piston relative to some fixed reference point such that increasing x corresponds to increasing volume in the cylinder. If you move the piston in or out by an infinitesimal amount dx, the volume decreases/increases by an amount dV = Adx, where A is the cross-sectional area of the piston. If you exert just barely enough force to get the piston going and keep it going at constant speed while you are doing your compression/expansion, the magnitude of the force, F you exert at a given instant of time will balance the force exerted by the gas on the piston. So, F = PA, where P is the pressure of the gas at that instant of time. The work done (by you) on the system in going from $x = x_i$ to $x = x_f$ is then

$$W = -\int_{x_i}^{x_f} PA \, dx = -\int_{V_i}^{V_f} PdV.$$

This formula is very important and I need to make a few comments at this point.

^{**} We shall explain the origins of this result later in the course. For now, read what the book has to say about this in §1.3.

• While we have derived our formula (work is integral of PdV) for a special case — cylinder with a piston — this formula turns out to be quite general. Whenever one *slowly* and steadily alters the volume of a system so that the external force per unit area balances the pressure exerted by the system and so that, in particular, the system always has a well-defined pressure and volume, one can use such a formula for the compressional part of the work.

• The minus sign in the work formula is important and easily forgotten. Don't. You can understand it as follows. Firstly, W is the work done by you on the system. With our definition of x, the pressure of the gas tries to increase x. So the force \vec{F} you exert is toward decreasing x. Thus the integrand in the usual definition of work,

$$W = \int_{\vec{x}_1}^{\vec{x}_2} \vec{F} \cdot d\vec{x},$$

is in this case given by

$$\vec{F} \cdot d\vec{x} = -PA \, dx.$$

The minus sign also goes with the convention that W is the energy entering the system. If you compress the gas $(x_f < x_i)$ you do positive mechanical work on the gas, raising its energy. If you let the gas expand $x_f > x_i$ the gas does mechanical work on you, lowering its energy. (Remember that pressure is always positive.)

• The formula for compressional work W is, up to a sign, given by the area under the graph of pressure vs. volume. But the shape of this graph – and hence the area under it – depends upon the way in which the pressure varies with volume. The dependence of pressure on volume depends on the details of the process whereby we change the state of the system. Whenever you need to compute compressional work, you will need to decide how P varies with V for the situation at hand. Mathematically, we need to figure out the function P = P(V) so we can do the integral. We will look at some important examples where we can figure out P = P(V) in just a moment.

• This method of computing the work is really only an approximation which becomes increasingly valid the more slowly we move the piston. The reason is that by moving the piston we disturb the equilibrium of the gas and, strictly speaking, observables like pressure are not so well-defined during the process of changing the state. Indeed, if you move the piston too rapidly the gas molecules "bunch up" near the cylinder so that there is no single observable pressure for the gas. But if we move the piston sufficiently slowly (speed of piston is much smaller than speed of sound in the gas, as it turns out) this problem becomes negligible and the integral expression – where we idealize the process as a sequence of equilibrium states – becomes a very good representation of what is going on during this process. Such an idealization of the process is called *quasi-static* and we shall discuss such things in a little more detail later. For now, just remember that our

compressional work formulas and things derived from it will really only be 100% accurate in the quasi-static approximation.

We now will have a look at some important processes where P(V) and hence the (quasi-static) work can be determined explicitly.

Example: Isobaric process

The simplest case for computing compressional work, mathematically speaking, would occur if the pressure were somehow kept constant as you move the piston. A constant pressure process is called an *isobaric* process. For and isobaric process we can easily do the integral to get

$$W = -P \int_{V_1}^{V_2} dV = P(V_1 - V_2),$$
 isobaric process at pressure P, ideal gas.

We have denoted the initial volume by V_1 and the final volume by V_2 . The isobaric process is, however, not that easily arranged with our piston scenario. (What do you have to do to get an isobaric compression/expansion?)

Example: Isothermal process

Another, pretty simple scenario – a doable one, too – arises if we manage to keep the temperature of the gas unchanged during the compression/expansion. For example, imagine that the cylinder is in thermal equilibrium with a large vat of liquid nitrogen and we do the compression/expansion very slowly so that the temperature change is negligible. Such a process is called *isothermal* — a code word which means T = const. Throughout an isothermal process at temperature T for an ideal gas we have P = NkT/V with T (and N) held fixed, so that

$$W = -NkT \int_{V_1}^{V_2} \frac{1}{V} dV = NkT \ln\left(\frac{V_1}{V_2}\right), \quad \text{isothermal process, ideal gas.}$$

Note that if V_1 is bigger than V_2 — we compressed the gas — the work is positive as it should be. Similarly, if $V_2 > V_1$ the work done is negative, which you should expect. (Do you remember what the graph of the logarithm looks like?)

For this isothermal process we can use the first law to determine the energy transferred to/from the gas as heat. This is one of the principal uses of the first law. Heat transfer will occur because the temperature is held fixed and for an ideal gas this means the energy is held fixed. As the gas is, say, compressed, energy is added to the system as work which tries to raise its temperature. The energy departs as heat and, ideally, if the process is sufficiently slow the temperature change is negligible. Mathematically, recall that $U = \frac{f}{2}NkT$, so that $\Delta U = \frac{f}{2}Nk\Delta T = 0$ because T doesn't change. We know that $W \neq 0$ so, according to the first law, heat transfer must cancel energy transferred by work:

$$0 = \Delta U = Q + W \implies Q = -W = NkT \ln\left(\frac{V_2}{V_1}\right)$$
, ideal gas, isothermal.

You should be able to see from this expression that for compression (expansion) energy flows out (in) as heat. Can you see it? Does this seem physically reasonable to you?

Example: Adiabatic process

If the cylinder with gas is insulated so that no transfer of energy via heat can occur, we say that the compression/expansion process is *adiabatic* — a code word which means Q = 0. There are now a couple of basic physical questions we would like to address. How much work is done for a given change in volume? For a given change in volume, is the work in the adiabatic case more or is it less than the work done in the isothermal case? Here we can use the first law to help us out. The work, whatever its value is, will cause the energy – and hence the temperature – to change according to

$$\Delta U = Nk\Delta T = W.$$

This result lets us immediately answer the second question about whether the isothermal or adiabatic process uses more/less work. For an ideal gas, an isothermal process keeps the quantity PV constant, *i.e.*, $P \propto 1/V$, while for an adiabatic process PV is not constant, but satisfies $\Delta(PV) \propto W$. Now, think about the area under the pressure vs. volume curve. For an isothermal process we have P = (const.)/V where the constant depends upon the temperature. For an adiabatic process the temperature changes (proportionally to the work done). Suppose, for example, we compress the gas. In the adiabatic case the work is positive, the energy increases and, hence, the temperature goes up. From the ideal gas law this means the product of pressure and volume is increasing during the adiabatic compression. If you sketch a graph of the two processes in the P-V plane you will see that the P vs. V curve for an adiabatic process will intersect two of the $P \sim 1/V$ isotherm curves so that the adiabatic P vs. V curve is steeper than the isothermal one passing through the same initial (P, V) value. The area under the curve – and hence the work – is greater in the adiabatic case. Roughly speaking, for a given change in volume you have to push harder on the piston for adiabatic compression compared to isothermal compression.

How to compute W? If someone told you the change in temperature you could work out ΔU and hence W from the first law, of course. But can we figure it out just from the ingredients given? That is, if someone just specifies the initial state and the change in volume, can we determine the work done? We can, but it takes a few steps. Combining the formula for W with the first law we easily derive the infinitesimal relation:

$$\frac{f}{2}NkdT = -PdV.$$

Eliminating the pressure using the ideal gas law and rearranging the result gives:

$$\frac{f}{2}\frac{1}{T}dT = -\frac{1}{V}dV.$$

Integrating both sides and rearranging we get

$$\left(\frac{T_2}{T_1}\right)^{f/2} = \frac{V_1}{V_2}.$$

If the initial state of the system is given, then we know the initial values V_1 and T_1 . Given the final value V_2 we can then solve for the final value T_2 :

$$T_2 = \left(\frac{V_1}{V_2}\right)^{2/f} T_1$$

From the initial and final values of T we can determine ΔU and hence W:

$$W = \Delta U = \frac{f}{2}Nk\Delta T = \frac{f}{2}Nk\left\{\left(\frac{V_1}{V_2}\right)^{2/f} - 1\right\}T_1.$$

Another useful fact about adiabatic compression/expansion can be obtained by using the ideal gas law in our relation between temperature and volume, above. We have that

$$PV^{\frac{f+2}{f}} = const.$$

See if you can derive this from the formulas above. You can now see how the pressure varies with volume and how this variation depends upon the composition of the gas,

$$P(V) = \frac{const.}{V^{\frac{f+2}{f}}}.$$

With P(V) in hand one can, again, compute the work done and recover the result above — a nice exercise.

Here is a practical application of the preceding analysis. Consider a cylinder in a car engine. What happens to the temperature of the air inside the cylinder when the piston compresses it? Well, it goes up, as you might expect. But by how much? How should we analyze this question? First of all, this process is neither isothermal nor adiabatic, strictly speaking. Does this mean we are out of luck? No! The adiabatic analysis above gives an excellent model of the process. Now, you will rightfully object that cylinders in engines are not insulated - heat transfer can occur (as you may have noticed when you touched a hot engine). But even if a process is not strictly adiabatic, it may be well-approximated as adiabatic if it occurs quickly enough. If the piston is very suddenly moved from its initial to final position there is insufficient time for much energy to flow to/from the system and the adiabatic assumption is pretty good.* This is actually the case in a real-world engine.

Ok, let's try some numbers. Suppose we take a volume of air at room temperature in a cylinder/piston and suddenly reduce the volume by a factor of 20 – what is the final temperature? Air is composed of diatomic molecules which can be viewed as having 5 degrees of freedom (3 translational and 2 rotational if you view the atoms as point particles). We then have (using the ideal gas model, the quasi-static approximation, and the adiabatic approximation)

$$T_2 = (20)^{2/5} T_1$$

Assuming $T_1 = 300$ K – about room temperature – we get

$$T_2 = (300 \text{K})(20)^{2/5} \approx 1000 \text{K} \approx 1340 \text{F}.$$

Pretty hot! This compression heating is used for ignition in diesel engines.

Quasi-static processes

As I mentioned before, the observables and states in thermodynamics have meaning only in equilibrium. The principal goal of thermodynamics is to see how this or that interaction changes the states. Typically the change in state involves "relaxation" during which the state is not in equilibrium, of course. (Recall our example of the glass of cold water in a warm room.) However, if the change in state happens sufficiently "gently", (not too quickly, not too extreme) then the relaxation can be viewed as a continuous sequence of equilibrium states. This is really only an approximation in general, but we can often set things up so that the approximation holds to a sufficient accuracy that the idealization of continuous change of equilibrium states is a great description of what is happening. Processes which can be approximated this way are called *quasi-static*. For many processes of interest we can arrange, or imagine, that this quasi-static assumption holds. I must emphasize that the laws of thermodynamics do not require any assumption about quasi-staticity. Quasi-static processes are just easier to analyze using the laws of thermodynamics. The laws of thermodynamics can be stated in a way which does *not* need the quasi-static approximation. In particular, the first law,

$$\Delta U = Q + W,$$

^{*} You should worry if we are violating the quasi-static approximation when we move the piston this quickly. In principle we might, but as it happens we can move the piston fast enough for the adiabatic model to be valid, while still moving the piston slowly enough to preserve the quasi-static assumption. We'll say more on this a little later.

refers to the difference in energy between two states, regardless of whether the process which changed the state was quasi-static.

We have been tacitly using the quasi-static approximation, however, when we studied the infinitesimal change in observables due to work and heat. In particular, we have used the notion that an infinitesimal amount of compressional work, δW can be treated via

$$\delta W = -P(V)dV,$$

so that compressional work computed via

$$W = -\int_{V_i}^{V_f} P(V)dV,$$

presupposes the process in question is quasi-static. Thinking back to our cylinder/piston example, if you move the piston quickly enough the process is not quasi-static. In this case the work you do is actually greater than that computed with the quasi-static assumption.

Conclusions drawn from such infinitesimal analysis need not hold for processes which can't be modeled as quasi-static, strictly speaking. For example, when we adiabatically compress a gas, if we don't do it slowly enough then we can't use the above formula to compute the work. If we compress the gas sufficiently rapidly, for example, the pressure in the cylinder is not defined since the air near the piston is behaving differently than the air far from the piston.

How slow do we have to go to get the quasi-static assumption to be valid? It depends. When we disturb the system (*e.g.*, move the piston in the cylinder) there will be a relaxation time which depends upon the system under consideration. For the piston-cylinder example, the relaxation time is on the order of the time for sound to travel along the cylinder. As a rule of thumb, if the time scale for disturbing the system is much longer than the relaxation time, the quasi-static assumption is a good one. So, if we compress the gas in the cylinder by moving the piston much more slowly than the speed of sound, we are good to go. If we move the piston with a speed comparable to the speed of sound we won't get very accurate results (quantitatively speaking) by analyzing the system using the quasi-static assumption.

Bottom line: If you use the PdV formula for computing work, you are using the quasistatic approximation.