

Cooling by Expansion

Now we're going to change the subject and consider the techniques used to get really cold temperatures. Of course, the best way to learn about these techniques is to work in a lab where experiments in the micro-Kelvin to milli-Kelvin range are being done. One of the things you might learn, that you won't find in K&K is that cryostats have to be mechanically isolated from their surroundings. Any energy that is transmitted to the interior of the cryostat winds up as thermal energy that the cooling apparatus must remove from the system.

To get to really cold temperatures requires several stages. The first stage is usually the cooling of helium to just above its liquefaction temperature. This is followed by liquefaction and pumping on the helium to get to a few tenths of a Kelvin, the helium dilution refrigerator to get to milli-Kelvins, and adiabatic demagnetization to get to micro-Kelvin. These days, one can use laser cooling and evaporative cooling of trapped atoms to get to nano-Kelvins. (Recall the *Physics Today* articles on the Bose-Einstein condensate, handed out with lecture 17.) These techniques are not discussed in K&K!

So, how does one do cooling by expansion? The basic scheme involves a working substance (perhaps helium in the lab, a banned CFC in an old refrigerator?). The refrigerator operates in a cycle. The working substance (called the gas from now on) travels the following path. From the cold volume (where it picked up some thermal energy) it goes to a heat exchanger (where it picks up more thermal energy from the hot gas, and then to a compressor. The compressor compresses the gas, increasing its pressure and temperature. The hot, high temperature gas is allowed to cool in contact with the outside world (room temperature). The gas then flows through the heat exchanger where it gives up heat to the cold gas on the way to the compressor. Then it goes through an expansion device, where it expands, does work on the outside world, and cools. Finally, the cold gas is sent back to the cold volume where it extracts heat.

The expansion device is often a turbine which allows the gas to expand approximately isentropically. If a given amount of gas has energy U_1 , volume V_1 and pressure p_1 on the input side of the expansion device and U_2 , V_2 , and p_2 on the output side, then the total energy going in is $U_1 + p_1V_1 = H_1$. (The pressure p_1 pushes the volume V_1 through the device.) Similarly, the energy coming out is $U_2 + p_2V_2 = H_2$. Then the work done on the expansion device is


$$W = (U_1 + p_1V_1) - (U_2 + p_2V_2) = H_1 - H_2 = \frac{5}{2}N(\tau_1 - \tau_2),$$

where we have assumed an ideal monatomic gas so $U = 3N\tau/2$ and $pV = N\tau$ and the enthalpy is $H = 5N\tau/2$. Note that the heat exchanger operates at constant pressure, no mechanical work is done by the heat exchanger, and the energy transferred in the form of heat appears as the change in enthalpy of the gas.

In a typical refrigerator for operation with helium, the gas is not liquefied in this step. There may be several stages of expansion cooling required to get the helium close to its liquefaction temperature. In household refrigerators and air conditioners the gas may be allowed to condense into a liquid. In this case the expansion device is often a porous plug which acts as the expansion valve in a Joule-Thomson device to be described shortly.

Throttling Processes

Consider a device with two insulated chambers. The chambers are separated by an insulating partition, with a small hole. The chambers have pistons which can be used to adjust the pressure in each chamber. Start with some gas (or whatever) in the left-hand chamber with volume V_1 and pressure p_1 . Since the partition has a small hole



in it, some of the gas will get through the hole into the right-hand chamber. We adjust both pistons continuously, so the gas in the left hand chamber stays at pressure p_1 and the gas in the right-hand chamber stays at pressure $p_2 < p_1$. Since $p_2 < p_1$, the gas gradually flows from the left hand chamber to the right hand chamber. No heat is exchanged and no work is done by the hole or the partition. The only work done is the work done by the pistons. This is an irreversible process. Changing the forces on the pistons by an infinitesimal amount will not cause the process to go in the opposite direction. There need not be friction or viscosity at the hole; in fact we assume there isn't. There is a pressure drop at the hole simply because it is small and molecules only run into it (and get through to the other side) every so often. This process is called a *throttling process*. We also assume that the speeds of the pistons are slow enough that we can ignore any kinetic energy of the gas.

So, we start with some amount of gas containing internal energy U_1 , with pressure p_1 , occupying volume V_1 on the left-hand side of the partition. We end up with the same amount of gas containing internal energy U_2 , with pressure p_2 , occupying volume V_2 on the right-hand side of the partition. The work done on the gas by the left piston is $W_1 = +p_1V_1$ and the work done on the gas by the right piston is $W_2 = -p_2V_2$. As we've already mentioned, there is no heat transfer. Therefore the change in internal energy of the gas is equal to the work done on the gas,

$$U_2 - U_1 = W_1 + W_2 = p_1V_1 - p_2V_2 ,$$

or

$$U_2 + p_2V_2 = U_1 + p_1V_1 ,$$

or

$$H_2 = H_1 .$$

The single small hole in our partition is simply to illustrate the point. In practice, the small hole may be replaced by a partition containing many small holes, by a valve, by a porous plug, or even by a wad of cotton! What's needed is a way to restrict (throttle) the flow and produce a pressure drop without extracting heat or work.

Note that in the refrigerator discussed in the previous section, the pressure drop was not produced by a throttling process. Instead, there was a turbine which extracted the work. That's why that process was not a constant enthalpy process.

Note also that since a throttling process is irreversible, the state of the gas as a whole is not an equilibrium state (part of it's at p_1 and part of it's at p_2). So we cannot really plot a throttling process on a p - V diagram.

Some more comments on enthalpy:

Recall

$$dU = dQ - p dV$$

and in a constant volume (isochoric) process the change in energy is just the heat added and

$$U_f - U_i = \int_i^f C_V dT ,$$

For the enthalpy,

$$dH = dU + d(pV) = dQ + V dp ,$$

and in a constant pressure (isobaric) process the change in enthalpy is just the heat added and

$$H_f - H_i = \int_i^f C_p dT$$

In the case of an adiabatic (K&K would say isentropic) process

$$U_f - U_i = - \int_i^f p dV ,$$

while

$$H_f - H_i = \int_i^f V dp .$$

In the case of a free expansion (an irreversible process) in an insulated container, there is no work done and the energy is constant. In the case of a throttling process in an insulated container, we've just shown that the enthalpy is constant. For a monatomic ideal gas

$$U = \frac{3}{2} NkT ,$$

and

$$H = \frac{5}{2}NkT.$$

As another example of the use of enthalpy, consider a chemical reaction which occurs in a beaker open to the atmosphere. This is a constant pressure process, so the heat generated or consumed by the reaction represents a change in enthalpy of the system. Chemical handbooks list reaction enthalpies.

The Joule-Thomson Effect

When a gas is close to its liquefaction temperature, the molecules are necessarily close enough to each other that there is a significant contribution to the total energy from the interaction energy between molecules. This contribution is negative. If we let the gas expand, the interaction energy decreases in absolute value (that is it gets more positive), so work must be done against the attractive forces between molecules. If there is no other source for this work, it will come from the kinetic energy of the molecules of the gas and the gas will cool. How can we arrange for there to be no other source for the work? An extreme way would be to let the gas expand into a vacuum. This motivates the idea of an expansion valve which creates a pressure drop between two volumes of gas. In other words, the throttling process just discussed.

A *Joule-Thomson* apparatus contains a valve or other flow restriction device like a porous plug which allows for a throttling process. As we've just seen, the gas goes through the expansion valve at constant enthalpy. That is $H_1 = H_2$. For an ideal monatomic gas, $H = 5N\tau/2$, so there is no cooling of an ideal gas in a Joule-Thomson apparatus! Of course, we argued in the previous paragraph that it is the interaction energy between molecules that is responsible for liquefaction and an ideal gas has no interaction energy so maybe we need a better model than the ideal gas.

How about a van der Waals gas? You showed in problem 7 of homework 7 that for a monatomic van der Waals gas

$$H(\tau, V) = \frac{5}{2}N\tau + \frac{N^2}{V}(b\tau - 2a),$$

where a and b are the van der Waals constants we introduced earlier. The first term is the enthalpy of an ideal gas and the second term is a correction for the long range attraction and short range repulsion of the van der Waals interaction. Consider a temperature defined by

$$\tau_{\text{inv}} = \frac{2a}{b} = \frac{27}{4}\tau_c,$$

where τ_c is the critical temperature defined earlier. The van der Waals correction to the enthalpy changes sign at the inversion temperature, τ_{inv} . At temperatures above the

inversion temperature, the van der Waals correction is positive, and increasing the volume (expanding) at constant enthalpy means the temperature must rise. (Not what we want!) Below the inversion temperature, the correction is negative and expanding at constant enthalpy requires that the temperature decrease.

Note that the inversion temperature fell out of our van der Waals model. What happens if the van der Waals model is not a good description? It must be that all gases have inversion temperatures, as all gas molecules have negative interaction energies at close (but not too close) separations.

If one combines a heat exchanger with an expansion valve, then one has a *Linde cycle*. High pressure warm gas comes into the heat exchanger where it gives up some heat to the low pressure cool gas going the other way through the heat exchanger. We assume that the input gas and the output gas have the same temperature. After the high pressure gas leaves the heat exchanger, it goes through the expansion valve where part of the gas liquefies. The remainder exits to the heat exchanger. Of course, the input and output of the heat exchanger must be connected to something—perhaps one or more stages of an expansion cooling device. The whole process occurs at constant enthalpy as before. Suppose the fraction which liquefies is λ . Then

$$H_{\text{in}} = \lambda H_{\text{liq}} + (1 - \lambda) H_{\text{out}} .$$

The three enthalpies are the enthalpy of the gas at the input pressure and temperature, the enthalpy of the gas at the output pressure and the same temperature (by the heat exchanger condition on the temperatures) and the enthalpy of the liquid at the temperature at which its vapor pressure is the output pressure. One can solve for the liquefaction fraction,

$$\lambda = \frac{H_{\text{out}} - H_{\text{in}}}{H_{\text{out}} - H_{\text{liq}}} ,$$

so if the output enthalpy is greater than the input enthalpy, some liquid is formed. This will always be true if the Joule-Thomson expansion cools the gas. K&K figure 12.4 shows some calculated values of λ for various temperatures and pressures. Values around 20% are typical.

Cooling by Pumping

Once liquid helium has been obtained, pumping on the liquid helium cools the liquid by evaporation. In the steady state, the temperature of the helium is set by the vapor pressure maintained by the vacuum pump. Typical temperatures that can be reached are from a few tenths of a Kelvin to a few Kelvin. Table 12.2 in K&K lists helium vapor pressures versus temperature for both ^3He and ^4He .

The Helium Dilution Refrigerator

As we remarked last time, at low temperatures ^3He and ^4He become immiscible. As $\tau \rightarrow 0$, the equilibrium configuration becomes essentially pure ^3He floating on ^4He containing about 6% dissolved ^3He . With such an arrangement, the chemical potential for ^3He must be the same in both the pure ^3He phase and the ^4He phase.

Now suppose we have such an arrangement, but we replace the ^4He containing dissolved ^3He with pure ^4He . Then the ^3He “evaporates” into the ^4He until the “vapor pressure” of the ^3He in the ^4He reaches equilibrium (which happens at 6% concentration as $\tau \rightarrow 0$).

In this evaporation process, latent heat is used up, and the liquid phase ^3He cools. If we arrange to replenish the pure ^4He in a cyclic or continuous method, we can make a continuous operation refrigerator. K&K describe several methods that are used to replenish the ^4He including distillation.

The latent heat that’s released is just the difference in chemical potentials. The internal energy of the pure ^3He is proportional to the Fermi energy (presumably constant) plus a term proportional to τ^2 . (See the discussion in lecture 16 on the heat capacity of a cold fermi gas.) As the temperature goes down, the latent heat available goes down and the practical limit of a dilution refrigerator is about 10 mk.

Isentropic Demagnetization

This technique is usually called adiabatic demagnetization. We've already mentioned it in connection with our magnetic spin system that we discussed in lecture 4. Consider a paramagnetic system. We don't want ferromagnetic interactions, and we need intrinsic magnetic moments, so we don't want a diamagnetic system either. Suppose the system is in contact with a thermal bath (probably provided by liquid ^3He !) at temperature τ_1 and an external magnetic field B_1 is applied. We need $\mu B_1/\tau_1 > 1$. When this is satisfied, most of the magnetic moments are aligned and the entropy of the magnetic system is very low. In fact, the distribution of spins depends only on $\mu B/\tau$, which means the entropy can only depend on $\mu B/\tau$. Now disconnect the spin system from the thermal bath and decrease the field to B_2 without changing the distribution of spins (leaving the distribution of spins constant is the isentropic part). Then the entropy doesn't change. Since the entropy didn't change it must be that $\mu B/\tau$ remained constant while the field was lowered. So the temperature had to change (it's basically dialed in with the field). The final temperature satisfies

$$\tau_2 = \tau_1 \frac{B_2}{B_1}.$$

The temperature of the spin system can be lowered substantially by this technique. The spin system isn't completely isolated, it's still in weak thermal contact with the lattice vibrations (phonons) of whatever solid the spins are contained in. (We assume the solid is an insulator, so we don't need to worry about the electronic contribution to the heat capacity. At low enough temperatures, the energy that can be absorbed by the spin system can be comparable to the phonon energy content of the lattice. With an electron paramagnetic system, temperatures in the milli-Kelvin or fraction of a milli-Kelvin range are possible. (I think K&K figure 12.9 must have its temperature axis mislabeled. It must be the final temperature in mK, not K!) The same trick can be played with nuclear magnetic moments which are several orders of magnitude weaker than electron magnetic moments. Nuclear adiabatic demagnetization must start from temperatures in the milli-Kelvin range and can reach temperatures around the micro-Kelvin range.

You might notice that the final temperature is proportional to the final magnetic field. Why not let the final magnetic field be zero? Then the temperature of the spin system (before it comes to equilibrium with the lattice vibrations) would be zero. Wrong! The external field can be taken to zero, but there will still be some internal magnetic field due to the magnetic interactions and inhomogeneities on the atomic scale. So there is some effective final value of B that's not zero which limits the lowest temperature the spin system may reach.

Laser and Evaporative Cooling

As we've already remarked, more modern techniques for getting really cold involve atom traps and laser and evaporative cooling. See the list of articles in lecture 17.

To get a feel for laser cooling, consider an atom in the path of two lasers pointed in opposite directions (actually it's the same laser, but with two beams directed at the atom from opposite directions). This will suffice for "1D cooling." To get the other two dimensions cold, we need 4 more beams. The lasers are tuned just below a transition energy of the atom (from the ground state to some excited state). If the atom is at rest (very, very cold), it sees the laser beam as off resonance and nothing happens. If the atom moves towards one of the laser beams, that beam appears blue shifted to the atom and starts to be on resonance. If the atom absorbs some of the light it gets a kick opposite to its velocity, slows down and is cooled! What I've just described is way over-simplified. The coherent oppositely directed lasers form a standing wave. A proper treatment of the situation involves a quantum treatment of the atom and its electron states in the standing wave field of the laser. Nevertheless, the simplified treatment gives an idea of what's going on.

Evaporative cooling is just what it sounds like, the high energy atoms are allowed to leave the trap and what's left behind are the lower energy and cooler atoms. Various tricks are used to produce interactions among the atoms in order that some will get some excess energy and evaporate. The article on a Fermi gas cited in lecture 17 describes such a trick.