An introduction to cryogenics

Adam L. Woodcraft
Home page:http://woodcraft.lowtemp.org
SUPA, Institute for Astronomy, Edinburgh University, Blackford Hill, Edinburgh EH9 3HJ, UK
(Dated: October 25, 2007. Last major change: October 23, 2007)

Contents

1. INTRODUCTION

These notes aim to give a somewhat basic introduction to the hows (and occasionally whys) of cryogenics. Cryogenics is not a very well defined term; here I have taken it to mean systems that operate at temperatures of 77 K and below, with the emphasis on systems operating below 4.2 K and above 1 K. The first two temperatures are, of course, the boiling points of nitrogen and helium respectively.

A cryogenic experiment or system is normally dominated by the need to get something cold and keep it cold, with other elements of the design subservient to that. Generally, the need to operate at cryogenic temperatures makes even an otherwise simple experiment complicated, and the colder you need to go, the harder life generally gets. Since it is so time consuming to get things cold, the ability to experiment with new techniques is somewhat limited, and most low temperature labs cling to a set of tried-and-trusted methods. Often these differ from lab to lab, and in the most interesting cases are contradictory.

First, a few basic definitions. In these notes, I shall use the following words in the way defined here (these are not the only definitions)

- **cryostat**: the system that a low temperature experiment goes in
- **dewar**: something for storing and (possibly) transporting liquid cryogens

Helium comes in two isotopes; the common helium-4 (^{4}He), which is obtained from oil and gas wells (it starts life as alpha particles) and the much more exotic helium-3 (^{3}He) which does not occur in nature, but is produced by the decay of tritium.

There are several books to which you can turn to expand on these notes, and after some years of neglect, old books are being reprinted in new editions, and new books appearing. The following lists the ones that I refer to the most; there are other books which are also well worth reading.

- Ekin (Ref. [1]); this is a new book, and gives lots of useful, detailed, down-to-earth advice on many aspects of cryogenics, though it is generally limited to temperatures above 1 K, and has an emphasis on making measurements on superconductors (you may or may not think this is a good thing).
- White and Meeson (Ref. [2]) is also a very worthwhile read, and is up to date.
- Pobell (Ref. [3]) is another standard textbook, and a new edition has recently appeared.
- Richardson and Smith (Ref. [4]) is a collection of seminars from the Low Temperature Group at Cornell University, and although somewhat out of date in places gives a lot of hands-on advice.
- Wigley (Ref. [5]) gives useful information on mechanical properties of materials.

2. PROVIDING A CRYOGENIC ENVIRONMENT

A. Cooling methods

Every cryogenic experiment needs some way of reaching low temperatures. There are two ways of doing this; using cryogens or using mechanical coolers.

Liquid cryogen cooled cryostats use the same principle as the ancient technique of using blocks of ice to keep food cold in storage and transport – incoming heat melts the ice, maintaining a constant temperature until no ice remains. These days, of course, we have replaced ice with mechanical refrigeration. With a liquid cryogen, evaporation of the liquid maintains the temperature at the boiling point, and this is how the majority of cryogenic experiments are cooled and kept cold, though mechanical cryogenic coolers have become more popular in recent years.

1. *Liquid cryogens*

a. Cooling with cryogens. Cryogens provide cooling in two ways:

– Heat is used in **evaporating** the liquid - it absorbs **latent heat**.
– Heat is then used in warming the gas from the boiling point to room temperature; the heat required to do this is given by the **enthalpy** change.

The two main cryogens in use in cryogenics are helium and nitrogen; some properties are given below (B.P. is the boiling point at 1 atmosphere, and the costs are of course approximate):

Cryogen	B.P.	Latent heat	Cost/litre
Nitrogen	77 K	161 kJ l^{-1}	£0.1
Helium-4	4.2 K	2.6 kJ l^{-1}	£3

If you are only going to cool an experiment to 77 K, then liquid nitrogen is the sensible solution. If you need to cool your system down to 4 K, then it makes sense to first cool to 77 K with liquid nitrogen, and then introduce helium, since nitrogen is about a thousand times cheaper than helium per Joule of latent heat. This works out particularly well because the heat capacity of components in a cryostat drops sharply with temperature, so by the time it has reached 77 K, most of the heat has already been removed from the cryostat.

We can see how useful this approach is from the following figures, which compares the quantity of cryogens required to cool 1 kg of copper [3] with and without the use of nitrogen to pre-cool the system:

Cryogen	Temperature range	Amount evaporated
Helium-4	300 K to 4.2 K	32 l
Nitrogen	300 K to 77 K	0.46 l
Helium-4	77 K to 4.2 K	2.2 l

While the latent heat of helium is rather small, the enthalpy between 4.2 K and 300 K is much larger; about 200 kJ/l. We can therefore reduce our consumption of cryogens even further by using the enthalpy. If all the enthalpy were utilised, we would have the following consumption [3]:

Cryogen	Temperature range	Amount evaporated
Helium-4	300 K to 4.2 K	0.8 l
Nitrogen	300 K to 77 K	0.28 l
Helium-4	77 K to 4.2 K	0.16 l

We can see that while the saving in nitrogen is not spectacular, if we can use all the enthalpy, we can save a large amount of helium. In practise, this is not possible. While you are guaranteed to use all the latent heat of a cryogen (it will all boil away eventually), there is no such guarantee for the enthalpy; the gas must leave the cryostat at room temperature in order for all the enthalpy to be used. This is hard to achieve. However, you can certainly arrange to maximise the gain from enthalpy. For example, you can design your cryostat so that the gas is slowed down as it leaves the cryostat, by passing through baffles or a medium such as cotton wool.

Helium storage dewars generally have a set of radiation shields which are thermally anchored to the neck, enabling them to be cooled by the enthalpy of the boil-off gas. Some cryostats have a liquid nitrogen bath used to cool radiation shields to 77 K to reduce the amount of heat reaching the 4 K stage. This is particularly common when optical access is required (for example for an astronomical instrument) and there is a filter to be cooled at 77 K.

Finally, it is important to note that nitrogen has a rather large latent heat of melting[1]. It is therefore very important to remove all the nitrogen before starting to fill with helium, otherwise you will waste a large amount of helium to solidify the nitrogen.

Hydrogen and neon are occasionally used as cryogens (boiling points at 1 atmosphere of 20 K and 27 K respectively). However, neon is ten times as expensive as helium, and hydrogen is explosive when mixed with air.

b. Handling, storing and transferring cryogens. Cryogens are usually stored in a transport dewar, either filled locally (from a nitrogen tank or helium liquifier), or delivered along with their contents (which will then boil off at one or two litres/day). Transferring nitrogen is fairly straightforward; it can be transferred using flexible tubing, or poured (carefully) from one container to another. The high latent heat means that a polystyrene bucket can hold nitrogen for several hours.

Transferring helium is much harder work. The low latent heat means that even transfer tubes must be vacuum jacketed. These vary in complexity from simple rigid affairs to long flexible tubes (much more expensive, but well worth the money). The basic principle is that a transfer tube is sealed into the storage dewar and the cryostat, and the storage dewar is pressurised (either using helium gas from a cylinder or by squeezing a bladder attached to the dewar), thus causing helium to be transferred. In practise getting a transfer right can be tricky, with some cryostats being much easier than others.

As discussed above, pre-cooling with liquid nitrogen is almost universal. Make sure you get all the nitrogen out after pre-cooling; this is usually done by pressurising the bath and forcing it out through a tube which reaches the bottom; make sure it really *does* reach the bottom! To get the most from the enthalpy, it's usually best to transfer helium into the bottom of the cryostat for the initial cool-down, so that the cold gas passes up past the rest of the cryostat, providing cooling. This is particularly important when there is a magnet to be cooled down. If necessary the cryostat should have some system for permitting this. Subsequent transfers should be done into the top of the helium reservoir, otherwise warm gas pass-

---

[1] More accurately known as the latent heat of fusion

ing through the transfer tube at the start will boil off liquid helium that is still present.

A very good set of instructions is given in reference [1], but the best way to learn is to do the first few transfers with someone who knows what they are doing until you (and they) are confident that you know what you're doing. It is disturbingly easy to use up a whole dewar without getting any helium into your cryostat.

c. Safety. There are three main dangers associated with handling liquid cryogens:

- Their low temperature, which can cause severe skin damage in the same way as extreme heat. **Burns** can occur from contact with liquid cryogens, cold gas produced by evaporation of liquid cryogens, and with pipework containing cold gas (e.g. helium boil-off from a cryostat). In extreme cases frostbite may occur.
- The large difference between volume in the liquid and gas phase (about a factor of 700) which can lead to an **explosion** if cryogenic liquids are unable to escape from a volume.
- **Asphyxiation** due to displacement of oxygen by evaporated cryogens (see point above).

Therefore, when handling cryogens, use the appropriate protective equipment (gloves, face shield, etc.) to prevent burns; you should bear in mind that skin can stick to cold metals. Eye protection from cold gases is particularly important. Another good reason for eye protection is that some solids cooled to cryogenic temperatures become brittle, a common example being flexible tubing used to carry liquid nitrogen, which can shatter at unexpected moments.

The most serious accidents generally occur due to the explosive bursting of a dewar when evaporated gases have nowhere to go, often caused by build up of ice and frozen air in the outlet. Cryostats should all have safety valves designed to prevent this from happening, but nevertheless caution is required; make sure dewars and cryostats have a small positive pressure to avoid ingress of air *(and don't leave them open to the air for any longer than necessary)*, keep a close eye on them, and make sure the correct valves are open. A burst cryostat is an impressive sight but one best viewed after the occasion, and ideally when it belongs to someone else.

Precautions against asphyxiation include making sure that labs are well ventilated, routing exhaust gases outside the building, and having an oxygen depletion alarm in the lab. For obvious reasons, don't travel with a nitrogen or helium dewar in a lift (and make sure that nobody else gets into the lift with it while you're running up or down the stairs.)

Another fact to bear in mind is that nitrogen will evaporate first from liquid air, leaving an oxygen-enriched liquid and a potential **combustion hazard**. Precautions against this include keeping the lid on buckets of nitrogen, and keeping cold surfaces which may condense air (such as pipes carrying helium boil-off gas) clean (no grease, oil, or anything else that is readily combustible).

The above is not intended as a comprehensive guide to safety hazards with cryogens; your lab should have a code of practise for handling liquid cryogens – use it!

2. *Mechanical*

a. Cooling with mechanical refrigerators An alternative to using liquid cryogens to cool an experiment is to use a mechanical cooler. There are two types generally in use:

- Gifford-McMahon (**GM**) coolers
- Pulse tube refrigerators (**PTR**s)

Both consist of two components, linked by plumbing: a room temperature compressor (which may be water cooled), and an expansion chamber inside the cryostat. Both work on the principle of compressing helium gas at room temperature (or a little above) in the compressor, then sending it into the cryostat where it is expanded to produce cooling. The gas passes through a high heat capacity "regenerator" which cools the gas on its way into the cryostat, and is in turn cooled by the expanded gas as it leaves.

A G-M cooler uses a piston in the expansion chamber, whereas a pulse tube refrigerator avoids the use of cold moving parts by using acoustic pulses in the helium gas to replace the moving piston. This has the advantage of reduced vibration, and reduced maintenance costs.

Both types of cooler are commercially available, and can pay for themselves in a year or so by removing the need to buy liquid helium (they are incredibly inefficient, however, and use a large amount of electrical power; for example 11 kW to produce 1.5 W of cooling at 4.2 K). Typical cooling powers at 4 K vary from 0.5 to 1.5 W. Of course, avoiding liquid cryogens has other advantages; it makes operation much safer and more straightforward, and for labs without a liquefier avoids the need to plan ahead by ordering sufficient helium, and possibly ending up with it all boiling away if a cool-down has to be aborted. It also avoids being tied to a cryostat by having to be available to do helium transfers every couple of days... To cool to 4-K with a mechanical cooler requires a two stage system, with the first stage operating at a temperature of 40 to 80 K. The first stage typically has a large cooling power (e.g. 40 W at 45 K, compared with 1.5 W at 4 K), and can thus be used to cool radiation shields in a similar manner to a helium dewar with a liquid nitrogen bath.

The main downside to mechanical coolers (apart from high initial cost, and the possibility of repair bills) is the vibrations they cause. This can interfere with or even prevent experiments from being carried out, although with careful design the effects of vibration can often be mitigated. If you buy a cooler from the USA, it will generally be optimised to operate from a 60 Hz electricity supply, and will have worse performance when operated in Europe with a 50 Hz supply (moral: make sure they quote you the performance at 50 Hz).

Mechanical coolers are particularly attractive when cryostats are run in remote locations, such as on an astronomical instrument sitting half way up a large telescope on the top of a mountain. Sometimes they can be used along with liquid cryogens, in applications where the cooler cannot reach a sufficiently low base temperature to be used alone, but instead reduces boil-off of the cryogens (for example a rather exotic system used in an X-ray astronomy satellite where a mechanical cooler reduced boil-off from a solid neon dewar, which itself reduced boil-off from a liquid helium dewar).

b. Safety Most of the points above on safety with liquid cryogens don't apply to systems using mechanical coolers, but some care is still required (for example it is still possible to burn yourself with a mechanically cooled cryostat by opening it up before it is warm).

3. *Cooling below 4 K*

So far I have discussed methods of cooling to 4 K. Cooling to lower temperatures involves further complexity, and I will give a very brief summary here.

a. Pumped helium-4 systems Helium-4 boils at 4.2 K at one atmosphere, but like any gas, the boiling point will reduce as the pressure is reduced. Therefore, by pumping on the liquid we can reach a temperature of about 1 K (below this the vapour pressure become too small). This can be done by pumping on the main bath of a liquid helium dewar, or, more commonly, on a small container (**1-K pot**) which is continually replenished with helium from the main bath. The latter scheme is popular since it means the cryostat can be re-filled without warming to 4 K, and is somewhat more efficient, since you don't waste cooling power cooling the helium in the main bath (helium has a very high heat capacity, and around half the liquid is lost when cooling from 4.2 K to 1 K).

Currently, commercial mechanical coolers will not reach 1 K. However, it is possible to use a self-contained fridge containing helium-4, in which the helium-4 is pumped using a **sorption pump**; this relies on the fact that charcoal adsorbs helium at low temperatures, but the helium can be driven off by heating the charcoal to above 20 K. In operation, the charcoal is heated and the helium driven into the 1 K pot. The charcoal is then allowed to cool, pumping the helium back out of the 1 K pot and providing cooling. The disadvantage of such a system is that it does not provide continuous cooling, although a system of two separate fridges can be used to provide quasi-continuous cooling. Commercial mechnically cooled systems are also available in which Joule-Thompson expansion of helium is used to reach approximately 1.5 K.

b. Pumped helium-3 systems Helium-3 maintains a useful vapour pressure to lower temperatures than helium-4, and pumping on helium-3 can produce temperatures lower than 300 mK. Helium-3 does not occur in nature and is consequently very expensive (more than 10 000 times as expensive as helium-4). It must therefore be handled in a closed system, unlike helium-4 which is commonly pumped away into the atmosphere. **Helium-3 refrigerators** can be constructed using an external gas handling system to pump on the helium-3 pot and return it to the cryostat, or they can be self contained using sorption pumps. It is usual to combine sorption systems with a helium-4 sorption refrigerator to produce a self-contained system which can be bolted onto the 4 K plate of a helium or mechanically cooled cryostat to produce temperatures below 300 mK, albeit with a rather low cooling power. Such systems are popular for cooling detector systems in astronomy, fusion diagnostics and other fields.

c. Dilution refrigerators Cooling below 300 mK requires more sophisticated techniques. The most common is a **dilution refrigerator**. These are commercially available from several companies in Europe, the US and Japan. I will not go into the details of how they operate here (copious information is available in many books (for example Refs. [2, 3]) other than to say that cooling is produced by mixing helium-3 and helium-4 at low temperatures. This can be used to achieve temperatures of a few mK.

To operate a dilution fridge requires a helium-3 circulation system similar to that used for a helium-3 fridge. Usually a 1-K pot is used to pre-cool the incoming helium-3. However, a number of companies now offer dilution refrigerators operating with a pulse tube cooler, rather than a liquid helium bath. These have to operate without a 1 K pot since there is no helium available to run one from. It turns out that this is possible with a change in design along with modifications to the gas handling system at room temperature, but it makes operation of the fridge somewhat more complicated. Recently, self-contained dilution refrigerators using sorption pumps have been developed.

d. ADRs An alternative to the use of a dilution fridge is to use magnetic cooling, with an **adiabatic demagnetisation refrigerator**, or **ADR**. The general principle is that a paramagnetic salt is magnetised in a field of, say, 6 Tesla. This will produce heating, which is taken away by contact with a helium bath[2] at a temperature of 4 K. This contact is arranged to be via a thermal link which can be opened (either by a solenoid or rod emerging from the cryostat). Reducing the field will then cause the salt pill to cool, to temperatures below 100 mK. Variations on this scheme are possible; to reduce the base temperature and improve performance they can be operated from an upper temperature of 1 K, or two or more salt pills can be operated in series. The main disadvantages are low cooling power and lack of continuous operation, although systems are currently in development that use multiple salt pills to achieve continuous operation. The magnetic fields required can also interfere with experiments.

However, they have several advantages over dilution fridges. Firstly, they are less prone to failing, as they do not have the intricate system of pumping lines and capillaries that are prone to plugging up, leaking, or both. At room temperature, all that is needed is a high current power supply, rather than a system of pipework, pumps and valves. Secondly, it is extremely easy to change temperature. With a dilution fridge, it is easy enough to warm a sample with a heater, but time-consuming to cool it down again. With an ADR, any temperature can be reached within a few seconds just by adjusting the magnetic field. For some experiments, this control is extremely useful. Although until recently popular only for astronomical

---

[2] Mechanical coolers can also be used, but are less ideal because they have a fixed cooling rate, and a heat load larger than this will cause them to warm up. A cryogen bath, on the other hand, has an essentially unlimited cooling rate without warming above the “base” temperature – until the cryogen runs out, of course.

instrumentation, commercial systems for materials characterisation have appeared in the last few years, some coupled with mechanical coolers to produce near-automatic operation from room temperature to below 100 mK.

Adiabatic demagnetisation of metal nuclei can be used to provide cooling to below 1 mK, but that is outside the scope of this section.

B. Types of cryostat

There are various types of cryostat that are used. There are two basic possibilities: the sample to be measured can be immersed in liquid helium, or it can be placed in an evacuated chamber (usually called the **vacuum can**). The simplest type of cryostat for both types of measurements is a small device which can be inserted into the neck of a helium transport dewar. However, most experiments require something more sophisticated.

The first cryostats used glass dewars. These are limited in size, and somewhat fragile, and are not in common use nowadays, having been replaced by metal dewars.

There are two common configurations. In perhaps the most popular one, the sample is located in a vacuum can surrounded by liquid helium. This ensures that the outside of the vacuum can is at 4 K (or lower), but has the disadvantage that a helium-tight seal must be made which works at low temperatures. This is usually achieved by making an O-ring out of indium (a very soft metal); lead (e.g. fuse wire) can also be used, but is harder to get right. Some people have successfully replaced indium O-rings with rings of Kapton®, a polymer. In vacuum can systems, an **exchange gas** is usually put into the vacuum can during cool-down in order to speed thermalisation of the internal components. Helium is commonly used, but this interferes with leak-checking if a leak from the surrounding helium is suspected. Using helium-3 avoids this problem, but is a somewhat expensive solution. Hydrogen can be used down to 20 K, but has to be handled with care because of the explosion risk when mixed with air, and of course nitrogen can be used as far as 77 K.

An alternative cryostat design has the experimental area outside the helium reservoir. (Of necessity, this approach is adopted with a cryocooled system). This has the advantage that the vacuum seal is made at room temperature, and can thus use a conventional rubber O ring. Such joints are less likely to fail, and are a lot easier and faster to make than a low temperature seal. The main disadvantage is that since the vacuum space reaches the outside of the dewar, it is not possible to speed up the cool-down of internal components by using an exchange gas.

3. GETTING AND KEEPING YOUR SAMPLE COLD

So far I have discussed how to produce and maintain a cryogenic environment. In an experiment, there will generally be something (I'll call it the "sample") that needs to be cooled. If it can be placed in the cryogenic liquid, then successful cooling is almost assured. In general, however, you will want to put it in vacuum, in which case making good thermal contact becomes important. In addition, the sample will be connected to the outside world, almost certainly by wiring, and possibly in other ways (such as via fill lines for a sample involving helium). These connections are of course routes for unwanted heat to enter the sample, and this has to be taken into consideration. While heat transfer is usually the most important consideration, heat capacity and thermal expansion can also be important. Getting all these things right involves a knowledge of material properties at low temperatures.

The good news is that there are many places in which material property values can be looked up, including the popular cryogenics text books, some of which I listed in the introduction, and in particular Ref. [1], which has an extensive selection of tables and graphs. The bad news is that there is no single comprehensive source of information, so for a relatively obscure material you may have to search in several places, and in the end go looking for individual papers. Ref. [6] lists of lot of sources of information. The closest thing to single comprehensive source is the set of books usually known as Touloukian (e.g. Ref. [7]), but they are over 35 years out of date, and not very good for temperatures below 1 K[3]. Even worse, especially for thermal conductivity, there has been little attempt to critically analyse measurements from the literature to produce recommended values.

There is a rather limited selection of information on the world wide web; see Refs. [9] and [10]. The latter should be used with some care, as for some materials, insufficient detail is given on what the material actually is[4]. Useful information on the general behaviour of material properties are given in various textbooks, in particular Refs. [2] and [3].

A. Thermal conductivity

The material property which is usually of greatest interest is thermal conductivity. Variation between materials at cryogenic temperatures is much greater than at room temperature, with ranges of many orders of magnitude. For materials with high conductivity, as the temperature is reduced from room temperature, the conductivity will increase, reach a peak, and then decrease. For other materials, the trend is generally downwards with decreasing temperature. Values for a few materials are shown in Fig. 1 as a guide to general behaviour; more detailed information can be found in the sources mentioned above.

There are two main obstacles to looking up the conductivity of a material:

---

[3] A slightly more up to date version will be found in Ref. [8], but this only covers elemental materials.

[4] And note that you have to get the capital letters in the right place in the URL!

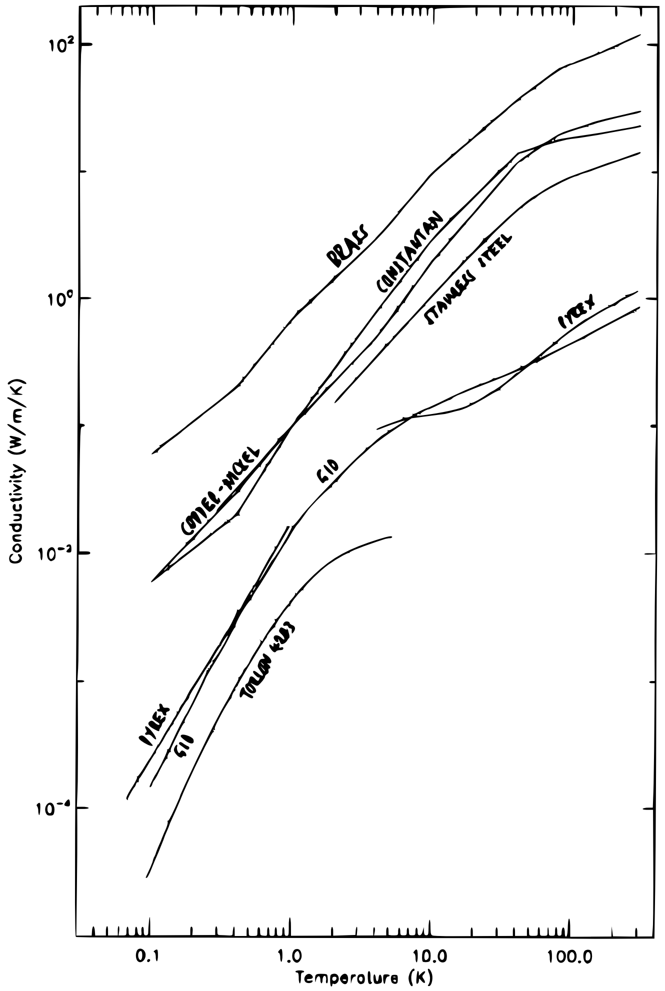

Figure 1: Approximate thermal conductivity values for some commonly used materials (brass, constantan, copper-nickel, stainless steel, Pyrex, G10 and Torlon 4203). Pure metals are not shown as their conductivity varies hugely with conductivity (see Fig. 2). From [3, 5, 11, 12].

– Many materials have just not been measured, or if they have the results were either not reported, or appeared in a sufficiently obscure place that you can't find it.
– There can be considerable variation between samples of nominally the same material, so the values you look up may not be correct for the material you have. (And sometimes, measurements in the literature are just wrong).

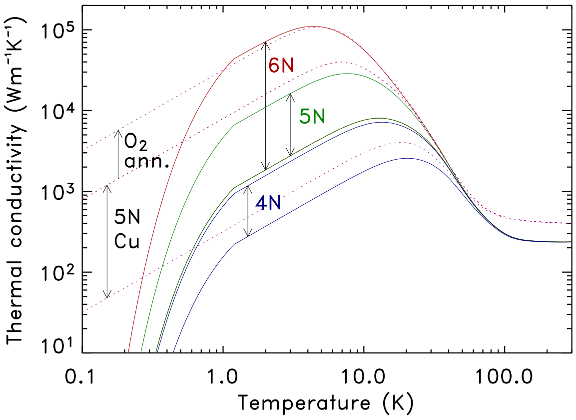

Figure 2: Recommended ranges for the conductivity of annealed aluminium with 4N, 5N and 6N purity (solid lines). Note that the lower limits for 5N and 6N purity are the same. Recommended values for vacuum annealed 5N purity copper are also shown, along with the higher conductivity values possible following oxygen annealing (dashed lines). From Ref. [13].

The latter point is particularly important in the case of pure metals, where the conductivity can vary by many orders of magnitude depending on purity. To take an example, the conductivity of copper at 1 K can vary from below 50 W/m/K to nearly 10^5 W/m/K depending on purity. It is true if one has some idea of the purity it is possible to reduce the likely range somewhat, but the uncertainty can still be considerable, as can be seen from Figure 2. One solution to this problem is to make use of the **Wiedemann-Franz** law, which relates the electrical and thermal conductivity of a metal:

$$\kappa = \frac{LT}{\rho}, \tag{1}$$

where L is the **Lorenz number**, and T is temperature. For sufficiently low temperatures, L is expected to take the theoretical value of $L_0 = 2.45 \times 10^{-8}$ WΩK^{-2}. Measuring the resistivity of a sample is relatively straightforward, especially if you have a helium transport dewar around to dip the sample into. However, very pure metals can require a rather large current in order to get a measurable voltage drop. Magnet power supplies are very useful here, as they are designed to put a large and well controlled current through a very low resistance.

For alloys, the resistance is much easier to measure, and a resistance bridge as used for reading resistance thermometers can be used. However, for low conductivity alloys the Wiedemann-Franz law will underestimate the thermal conductivity, because the heat transported by electrons becomes sufficiently small that conduction through the crystal lattice will also be significant, and this has no electrical counterpart.

Once you have a value for conductivity, $\kappa(T)$, calculating the heat leak P for a material with area A and length L, between temperature T_1 and T_2, is straightforward:

$$P = \frac{A}{L}\int_{T_1}^{T_2} \kappa(T)dT. \tag{2}$$

Often we can represent conductivity as a power-law:

$$\kappa(T) = \kappa_0 T^\beta, \tag{3}$$

and we can then write

$$P = \frac{A}{L}\frac{\kappa_0}{\beta+1}\left(T_2^{\beta+1} - T_1^{\beta+1}\right). \tag{4}$$

At sufficiently low temperatures (and 4 K is almost always sufficiently low), metals will have a value of $\beta = 1$. Crystalline dielectrics will have a value around $\beta = 3$, and amorphous materials (such as glass) have a more complex behaviour, generally having an approximate value of $\beta = 2$ below 1 K.

Finally, please remember the difference between *conductivity*, κ, measured in W/m/K, and an intrinsic property of a material, and *conductance*, G, measured in W/K, which is the property of a particular thermal link. The two are related by

$$G = \frac{A}{L}\kappa. \tag{5}$$

It is easy to remember which is which if you compare to resistance and resistivity, which people rarely mix up. (I don't know why conductivity and conductance are harder to get right, but they are, and I still get it wrong myself on occasion.)

B. Contact conductance

Usually, there will be one or more joints between your sample and the source of cooling. In the simplest case, the sample is directly bolted to the bottom of the helium bath, or the cold stage of a mechanical cooler. In more complex systems, there may be a system of thermal links stretching for a metre or more.

Making good thermal contact between two surfaces is not straightforward, and at very low temperatures the conductance across different types of joint can vary by over 9 orders of magnitude! The general principle for good thermal contact is to join clean surfaces with as high a pressure as possible. The most common material for thermal contact is copper, which forms an insulating oxide layer. This can be removed every time a joint is made by filing the surface, but plating with gold is a much more satisfactory solution. Pressure is normally applied by bolts or screws, and in general it is the total force that is applied, not the area of the joint, that matters. Well prepared joints can have a conductance that approaches that of the bulk material. If it is not possible to apply large forces, or to have clean contacts, a small amount of vacuum grease can help by increasing the contact area. However, although it can make a bad joint better, it is likely to make a good joint worse, and the lower the temperature, the less effective it will be (since the conductivity of the grease drops much more sharply with temperature than for a metal).

C. Heat capacity and thermal expansion

Heat capacity and thermal expansion are also important, though usually not as critical as thermal conductivity. Unlike conductivity, they are not very sensitive to small variations in material composition, so using values from the literature is more accurate.

The heat capacity of most materials is approximately described by the **Debye model**, which characterises them by the Debye temperature, Θ_D, which is around room temperature for many metals. Above the Debye temperature, heat capacity approaches a temperature-independent value. Well below this (say $T < \Theta_D/25$), the heat capacity varies as T^3. This is not the full story, since this only describes the heat capacity due to lattice vibrations. Electrons will also contribute a linear component, though this usually only becomes significant at temperatures below 10 K. It is also important to know that some non-metals, in particular those with an amorphous structure, have very high heat capacities at low temperatures.

The thermal expansion coefficient reduces with decreasing temperature, and very little takes place below 77 K. Furthermore, the temperature variation of thermal expansion is similar for almost all materials. Reference [14] gives a critical analysis of thermal expansion data from the literature for many materials, along with recommended values.

D. Thermal radiation

Everything emits thermal radiation. A body with **emissivity** ϵ and temperature T will emit energy at a rate of

$$E = \epsilon\sigma T^4, \tag{6}$$

where $\sigma = 5.67 \times 10^{-8}\mathrm{Wm}^{-2}\mathrm{K}^{-4}$ (Stephan's constant). A blackbody corresponds to $\epsilon = 1$.
To reduce heat on our sample from thermal radiation, we can do two things:

– Reduce the temperature of the surroundings.
– Reduce the emissivity of the surroundings.

We achieve both of the above by surrounding the sample with a **radiation shield** at as low a temperature as practical, and by reducing the emissivity of the radiation shield as much as we can. Some emissivity figures are given in the table [2, 15]:

Material	Emissivity
Polished copper	0.02-0.04
Lightly oxidised copper	0.1
Highly oxidised copper	0.6
Aluminium (pure and alloys)	0.01-0.06
Highly oxidised aluminium	0.2
Gold	0.015-0.03
Ice[a]	0.9

[a] Greater than about 10μm in thickness

From this we can see that a polished copper or aluminium shield would do nicely. However, such a shield will not stay polished, particularly if handled, and in practise a gold plated copper shield is the best solution. (Aluminium is a bit tricky to make good thermal contact to, and is best avoided for shields unless it is really necessary to save weight[5]).

It should be noted that the peak wavelength of thermal radiation changes with temperature, and thus measurements of

---

[5] It also superconducts below around 1 K, depending on alloy type, resulting in a poor thermal conductivity.

emissivity at room temperature should be used with some caution to apply to cryogenic systems.

E. Other heat transfer mechanisms

1. Gas conduction

Conduction through gases can be a blessing and a curse. On the one hand, introducing an exchange gas is a very useful way of getting your experiment cold in the first place, when different parts of the experiment will operate at different temperatures and you have therefore carefully designed a poor thermal contact between them. On the other, even a small leak can add sufficient helium gas to what is supposed to be a vacuum space to stop your experiment from working.

The conductivity through a gas does not depend on pressure, unless the pressure is so low that it enters the ballistic regime, where the mean free path becomes larger than the size of the container. Therefore there is no need to get carried away with adding exchange gas - a pressure of around 0.1 mbar at 300 K is sufficient, and no matter how impatient you are getting, adding more gas isn't going to speed things up.

2. Heat switches

Sometimes it is not possible to use an exchange gas. There are other tricks for making thermal contact between parts of the experiment that must be thermally isolated when cold. You can use a **heat switch**, in which thermal contact can be turned on and off. This can be done with a movable mechanical contact, or with a tube full of gas which freezes out at the operating temperature. Other tricks include materials which go superconducting at the operating temperature (superconductors are very good conductors of electricity but not heat), sometimes referred to as a passive heat switch.

3. Film creep

Below 2.17 K, helium-4 goes superfluid. Superfluid helium (sometimes referred to as helium-II), has many weird and wonderful properties. One of them is that a superfluid film will creep up the sides of a container containing helium-II. At some point it reaches a temperature where it evaporates, being replaced by more helium flowing up through the film. This means that if you pump on superfluid helium, the film will try to escape up the pumping line. You stop this happening with a restriction in the pumping line.

4. SOME OTHER ASPECTS OF EXPERIMENTAL DESIGN

Obviously the details of cryogenic experiments vary widely, but there some aspects that are common to almost all low temperature experiments.

A. Wiring

There are two main ways you can go wrong with wiring. Firstly, you can end up dumping too much heat into your experiment. This might stop it from doing what you want altogether, or more subtly it may appear to work but you get the wrong result (for example, heat leaking down thermometer wires can cause the thermometer to read a higher temperature than whatever it is supposed to be measuring). You avoid this by proper **heat sinking**, also known as thermal anchoring, generally achieved by wrapping the wire several times round a copper bobbin. Depending on how cold your experiment is, you may need to heat sink at several different temperature stages.

The Wiedemann-Franz law says that electrical and thermal conductivity are proportional to each other, so in principle it doesn't matter what material you use for wiring; you just choose the cross-section to get an acceptable heat leak. In practise this would mean ludicrously small cross-sections if you used a pure metal such as copper, so alloys such as constantan and manganin are used. These also have the advantage that the change in resistivity as a function of temperature is much smaller than for a pure metal. At low temperatures, superconducting wire such as NbTi can be useful since it gets you round the Wiedemann-Franz law, offering high electrical conductivity and low thermal conductivity.

The second way you can go wrong is if the wiring goes open circuit when the experiment is cold. When this happens, the failure is usually at a soldered joint, though if you aren't careful thermal contraction could also cause a wire to snap. Since there is no way of getting inside a cold cryostat, if it was a vital wire, you have to warm up. Extreme care is therefore required when soldering wires in a cryogenic experiment, and using an eyeglass to inspect joints thoroughly and re-doing them if you have any doubt whatsoever is not being too paranoid.

B. Choosing materials

The mechanical properties of materials used in cryostats are important. Some metals go **brittle** when cooled down. You are safe with anything with an fcc (face centred cubic) structure, but in practise, low temperature labs usually stick to a small number of well known materials for the majority of uses. For metals, these are generally stainless steel, copper, brass and – sometimes – aluminium). There is a larger number of commonly used non-metals, including Teflon®, Mylar®, Kapton®, G10 (the green stuff that printed circuit boards are made from) and various epoxies.

Brass is a very useful material; it has a reasonable thermal conductivity and is easy to machine. On the other hand, it is not particularly well characterised, and sometimes recycled which occasionally results in the inclusion of unwanted materials. Folklore says that extruded bar can contain longitudinal cracks that could cause leaks, so anything that has to be leak-light is best constructed from plate.

Copper comes in many forms. Oxygen Free High Conductivity (**OFHC**) copper is popular in low temperature labs. OFHC copper is useful if it is to be hydrogen brazed, as oxygen reacts with the hydrogen resulting in porous joints. However, despite a common feeling that if it has "High Conductivity" in the name it must be good, the thermal (and electrical) conductivity is no better than regular commercial copper (electrolytic tough pitch), and possibly worse. Beware of copper tubing; although it is a tempting material to make radiation shields from, it is likely to be phosphorous-deoxydised copper, which has a very poor thermal conductivity.

Sometimes you want a good thermal but poor electrical conductor. Sapphire does this quite nicely around 4 K, but is not useful at ultra-low-temperatures. Since the conductivity of a dielectric varies with temperature as T^3, the conductivity drops[6] by a factor of over 60 000 between 4 K and 100 mK, whereas that of a metal drops by a mere factor of 40.

Another time where it pays to get the right material is the choice of tubing for transferring liquid nitrogen (and boil-off helium gas). Cord-reinforced PVC tubing is easily available and will do the job - for a while, but is rather prone to shattering when cold. Silicone rubber is the answer; readily and cheaply available (it's used in the food industry), and so long as you're careful not to stress it when it's cold, it rarely shatters.

C. Grounding and RF

An area that is very important, but hard to get right, is electrical grounding and the reduction of RF (radio frequency) interference. RF interference can affect your experiment, and in particular can cause thermometers to warm up and thus give incorrect readings. It can be reduced by filtering wires where they enter a cryostat, and in extreme cases putting the entire cryostat in an electrically shielded room. Further filtering at low temperatures can help. Getting grounding right is a complex task and one where it really pays to find an expert. If none are to hand, some advice on grounding (as applied to making the film "Jaws") can be found in Ref. [4].

5. WHEN THINGS GO WRONG

Many things can go wrong with a cryogenic experiment, and here are two of the more popular ones.

A. Leaks

Leaks can be the bane of the low temperature physicists life, especially with complex systems such as dilution fridges. If you are lucky(?) the leak is present at room temperature; it is much harder to find a leak if it only appears when the cryostat is cold. Unfortunately, this is all too common an occurrence. Even worse is one which only lets superfluid helium through...

The main weapon in the hunt for leaks is the **leak detector**; a mass spectrometer which is designed to sense helium (fortunately for us, helium is the standard gas used for leak detection in industry). Searching for a leak requires experience and it is a good idea to get someone who has done it before to help you. A leak detector with a graphical display is very useful for convincing yourself you really have found a leak.

There are various methods of leak-hunting, depending on where the leak is. If it is possible to pump on the system containing the leak (and the helium background inside it is not too high), then helium from a gas cylinder can be squirted around the various candidate locations for the leak through a narrow tube while the leak detector is used to pump on the offending system. Start at the top and work your way down, since helium is lighter than air. Sometimes it can be useful to increase the signal by putting a bag around a suspected leak and filling it with helium.

If the leak only appears when the cryostat is cold you have more of a problem. Usually (but not always) cooling to 77 K is sufficient to find the leak, in which case it is sometimes possible to find it by dunking the cryostat in liquid nitrogen, then pulling it out sufficiently to squirt helium round. This is not terribly gentle on the cryostat, though.

B. Blocked exhaust ports

A more urgent problem is a dewar or cryostat with an ice plug. This is another time where your first reaction should be to get experienced help. An ice plug can sometimes be freed by judicious prodding with a hollow metal pipe. If you do clear the blockage, you are likely to get a large stream of high pressure, cold gas, so wear appropriate protective equipment. It is a good idea to make sure you have the appropriate equipment ready before you cool down, as you don't want to be running round finding copper tubing etc. when you've just discovered that your cryostat has plugged.

Clearing a blockage is potentially dangerous, and if you have any doubts about what you're doing, don't do it. And it's much better to make sure you don't get a blockage in the first place.

6. WARMING UP!

Finally, at some point (perhaps days, weeks or even months) after you cooled your cryostat down, it's time to warm up. This is easy compared with everything else, but sometimes it can be worth putting heaters in suitable places to help speed things up. This is obviously something best thought of before you cool down.

And when you've got it warm, it is a good idea to make sure that everything has warmed above 0°C before cracking

---

[6] The moral of this is to beware of people with experience of cryogenics at 4 K "helping" you with your ultra-low-temperature experiment. It's a different world...

the vacuum and opening up, otherwise everything inside will cover itself with water frozen out of the air. If you *must* open it up early, then a heat gun can be used to drive off the moisture and finish warming everything up.

[1] Jack W. Ekin. *Experimental techniques for low-temperature measurements*. Oxford University Press, Oxford, 2006.

[2] G. K. White and P. J. Meeson. *Experimental techniques in low-temperature physics*. Oxford: Clarendon Press, 4th edition, 2002.

[3] Frank Pobell. *Matter and Methods at Low Temperatures*. Springer, 1992.

[4] Robert C. Richardson and Eric N. Smith. *Experimental Techniques in Condensed Matter Physics at Low Temperatures*. Addison-Wesley,, 1988.

[5] D. A. Wigley. *Materials for low-temperature use*. Oxford University Press, 1978.

[6] J. G. Weisand II, V. Flynn, E. Thompson, and R. P. Reed. A reference guide for cryogenic properties of materials. Technical Report SLAC-TN-03-023, SLAC, 2003. Available at www.slac.stanford.edu/cgi-wrap/getdoc/slac-tn-03-023.pdf.

[7] Y. S. Touloukian, P. E. Liley, and S. C. Saxena. *Thermophysical properties of matter. Vol. 1. Thermal conductivity: metallic elements and alloys*. Plenum, 1970.

[8] C. Y. Ho, R. W. Powell, and P. E. Liley. Thermal conductivity of the elements. *J. Phys. Chem. Ref. Data*, 3 Suppl. 1, 1974.

[9] http://researchmeasurements.com/about.html.

[10] http://cryogenics.nist.gov/MPropsMAY/material%20properties.htm.

[11] O. V. Lounasmaa. *Experimental principles and methods below 1K*. Academic, London, 1974.

[12] G. K. White. *Experimental techniques in low-temperature physics*. Clarendon Press, Oxford, 1979.

[13] A. L. Woodcraft. Recommended values for the thermal conductivity of aluminium of different purities in the cryogenic to room temperature range, and a comparison with copper. *Cryogenics*, 45(9):626–636, 2005.

[14] R. J. Corruccini and J. J. Gniewek. Thermal expansion of technical solids at low temperatures. *NBS monograph number 29*.

[15] Ed. B. Hands. *Cryogenic Engineering*. Academic Press, London, 1984.