Using Peltier devices

Peltier-effect devices represent a rival solid-state method for cooling small items. Richard Lines discusses how they work and explains how to apply them. His design is illustrated by a design example for cooling a ccd.

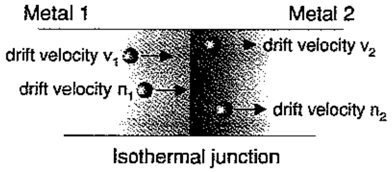

Due to Peltier effect, electrons crossing the dissimilar-metal junction energy change in energy level, producing or removing heat depending on the direction of current flow. With dissimilar metals, the effect is negligibly small, but when applied to a special semiconductor junction, Peltier effect becomes a useful means of extracting heat.

Thermoelectric coolers are convenient solid-state devices capable of pumping heat up a thermal gradient. They are very useful for cooling small items like charge-coupled device sensors, photodiodes and semiconductor lasers.

The principle of operation is closely related to the Seebeck effect, where a temperature difference between two junctions generates an emf. This effect was described in relation the thermocouple in my previous article.

The Peltier effect involves passing a current through a junction which produces a heating or cooling effect at the junction depending on the direction of the current.

The basic principle is outlined in **Fig. 1**. This illustrates electrons flowing from left to right across a metal junction. The solid state properties of the metal dictate the average energy of the conduction electrons – E_1 for metal 1 and E_2 for metal 2. Thus the electrons in metal 1 will bring energy up to the junction at a rate of,

$$E_{in} = n_1 E_1 v_1$$

and similarly energy will be carried away from the junction,

$$E_{out} = n_2 E_2 v_2$$

and the difference between these two terms is the energy taken from or given to the atoms in the lattice. Since the current must remain the same over the junction you can say that,

$$I = n_1 \times v_1 \times q = n_2 \times v_2 \times q$$

where q is the electronic charge, and I is the current.

Net heat flow from the junction is,

$$E_{out} - E_{in} = \frac{1}{q}(E_2 - E_1)$$

So the heat removed by a Peltier device is directly proportional to the current and the difference in energy carried by the conduction electrons.

The Peltier effect is due to the bulk properties of the two metals; in effect the heat capacity of the conduction electrons changes on crossing the junction. With normal metals, say copper/iron, the Peltier effect is very small and difficult to measure. In practice the normal ohmic heating effect is much larger. However, in semiconductors the effect is exaggerated and can be exploited as a useful method for cooling

A thermoelectric cooler consists of an array of p-n junctions wired in series between two thermally conductive ceramic plates. The p-n junctions are so wired that heat is absorbed on one plate, transferred through the semiconductor junctions, and leaves the device via the other plate; the device is a heat pump.

Heat is absorbed at the reverse biased junctions and emitted at the forward biased ones. A semiconductor material with a low bandgap is required – usually bismuth telluride – so that enough minority carriers can be generated at the operating temperature.

As the device is reversible it can also be used for heating, allowing an item to be maintained at constant temperature for ambient variations above or below the set point. For the sake of a consistent nomenclature, I will refer to the thermoelectric cooler surface connected to the controlled item as the cold plate and the surface connected to the heat sink as the hot plate, regardless of whether the thermoelectric cooler is cooling or heating the item.

The two main thermal parameters of interest for any thermoelectric cooler are the maximum temperature that can be maintained across the device and the amount of heat that can be pumped. The maximum temperature differential is quoted for no external heat load on the cold plate i.e. nothing being cooled.

Under these conditions all the electrical power input to the thermoelectric cooler is being used to overcome its internal reverse heat leak and heat input by radiation and convection through the area of the cold plate. This is the least efficient situation since virtually by definition the thermoelectric cooler is doing no useful work.

The maximum temperature differential figure ΔT_{max} is usually quoted for a given hot plate temperature in both a vacuum and dry nitrogen; running in a vacuum eliminates heat input by convection and allows slightly colder temperatures to be achieved. For the single-stage thermoelectric cooler as illustrated in **Fig. 2** the numbers would be typically 60 and

70°C for use in nitrogen and *in vacuo* respectively.

For applications needing a larger temperature drop, there are multistage devices available looking like small pyramids; the hot plate for the top layer is the cold plate for the layer underneath. A law of diminishing returns sets in; a three-stage device might have a maximum differential of 100º C but a six stage device may only manage 130º C.

How much heat?

The second parameter is the heat that can be input to the cold plate to reduce the temperature differential to zero, known as Q_{max}. Under these conditions all the electrical power is being used simply to prevent the cold plate getting warmer than the hot plate.

Since the cold plate temperature is the same as the surroundings there is no heat input due to convection or radiation and the thermoelectric cooler will be working efficiently. The maximum heat input scales with the area of the cold plate and is typically 2-3W/cm^2. There are tiny devices consisting of a single p-n junction which will pull only a few hundred milliwatts, up to larger 6cm^2 units rated at 100W.

Specifying a thermoelectric cooler

In order to select a suitable thermoelectric cooler it is necessary to determine how much heat is to be transferred over the required temperature difference. Two things really need to be known from the start. One is the cold plate temperature, which is decided by the application. The second is the method for getting the heat away from the hot plate, since this will define the hot plate temperature and thus the temperature drop across the thermoelectric cooler.

The second point usually boils down to the choice between a heat sink and a water loop. The water loop is more effective since the hot plate needs to be only a fraction of a degree warmer than the water even for quite modest flow rates to enable the heat to be disposed of.

A heat sink has to run warmer than the surrounding air to function, so the thermoelectric cooler has to work harder to achieve the same cold plate temperature. However the heat sink is more practicable in portable applications.

An example

As an example, consider a charge-coupled device image sensor to be cooled to –20ºC (253K) from room temperature with a heat sink, **Fig. 3**. For room temperature, assume 20ºC.

The heat sink is to run 10ºC warmer than the surrounding air, so the temperature drop across the thermoelectric cooler will be 50ºC. The ccd has case dimensions 20 by 15 by 5mm and exactly fits on the cold plate. It has 20 pins, all of which are connected to circuitry at room temperature by copper wires 0.25mm diameter and 100mm long.

Both ccd and thermoelectric cooler assembly are contained in dry nitrogen. Further, the ccd runs from a 10V supply with a current consumption of 5mA.

You are now in a position to estimate the total heat input to the cold plate. There are several terms involved. First is the active heat generated by the ccd due to its power consumption. In this case P_{active} is simply the voltage×current product, or 10V by 5mA=50mW.

As an aside, the temperature difference between the actual silicon slice and the cold plate is assumed to be zero. This means that that the thermal resistance, junction to case, is assumed very low. Usually, this is true, but there may be applications where the device dissipation changes in use, causing uncontrollable temperature changes in the silicon.

Remember that the sensor usually has to go outside the chip package! There are some photodiodes and ccds available which include a thermistor sensor bonded internally to the silicon chip to get around this problem.

There are now the passive components to be considered. These are the radiation, convection and conduction terms. You will need to know the exposed surface area of the ccd.

area=20×15=300mm^2 (front face)
+20×5×2=200mm^2 (sides)
+15×5×2= 150mm^2 (ends)

Total exposed area, A is 650mm^2.

Radiation. The exposed surface area is treated as a black body which will absorb radiation emanating from the surroundings at room temperature. This term is calculated using the Stefan/Boltzmann radiation law;

$$P_{rad} = A\sigma\left(T_{hot}^4 - T_{cold}^4\right)$$

where A is the area already calculated, T_{hot} is the background temperature assumed as 293K and T_{cold} is the ccd temperature at 253K. Symbol ρ is Stefan's constant, which is,

5.67×10–8W/m^2/K^4.

Absorbed radiation is found to be 0.12W. This is a worst case result since the surface will not behave perfectly as a black body, but it is obviously better to overestimate the heat absorbed.

Convection. Unless the unit is operated in a vacuum there will be circulating air currents inside the assembly'.The convection component is given by,

$$P_{conv}=Ah(T_{hot}-T_{cold})$$

where A is the exposed area, h is the convection coefficient is 21.7W/m^2/ºC. Temperatures T_{hot} and T_{cold} can be either ºC or K. The convection term is found to be 560mW.

Convection is often the largest contribution. It can be removed by operation in a vacuum; this also solves the problem of frosting up as the assembly is cooled past the dew point.

Conduction. The connecting wires to the ccd and temperature sensor form a heat leak which must be accounted for. Any retaining structure used to hold the ccd in place will also need to be considered.

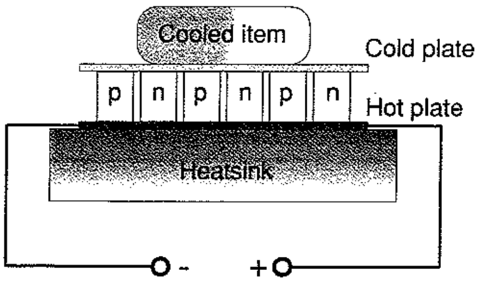

Fig. 2. When used for cooling, a Peltier device produces heat that usually needs to be removed via a heat sink.

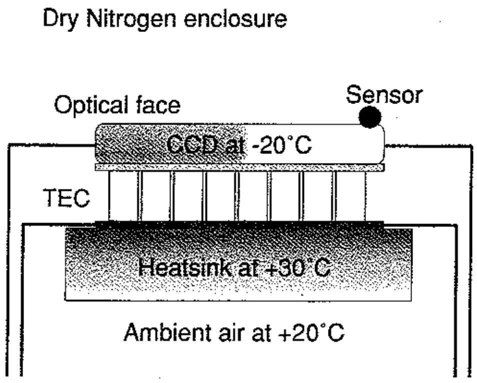

Fig. 3. Using a Peltier-effect device to cool a ccd.

Allowing three sensor wires, there will be 23 copper wires going to the ccd assembly. The equation describing heat flow for conduction is,

$$P_{cond} = \frac{ka}{l}\left(T_{hot} - T_{cold}\right)$$

where k is the thermal conductivity of copper is 386W/m/ºC, a is the area and l is the wire length, at 10cm. In this case, area a is 23 times the cross sectional area of a 0.25mm diameter wire.

Putting in the numbers shows that the conduction heat component will be 175mW. Note that using wire of twice the diameter will increase this heat leak by a factor of four.

Thus the total heat entering the cold plate will be $P_{active}+P_{rad}+P_{conv}+P_{cond}$ giving a grand total of 0.905W. This number should be treated as an approximation as all the terms have considerable errors. The active load could be in error by a factor of two either way since IC power consumptions are not that well defined as a rule. For the radiation term the emissivity has been assumed to be unity; the convection depends to some extent on the shape, angle and state of the surfaces.

Selecting a cooler from the range

Knowing the heat input, you are now in a position to consult the manufacturers' data sheets and choose a suitable thermoelectric cooler.

The information is normally presented graphically by a set of curves. These relate the current flowing through the thermoelectric cooler and the heat input with the temperature differential produced.

Figure 4 shows a typical set of curves. It so happens that the shape is very much the same for all single-stage thermoelectric coolers so the one set can be used for many devices simply by normalising the axes to suit.

The *y* axis is the temperature drop produced typically normalised to 60ºC, so the 50ºC needed in our application will be represented as a horizontal line at 0.83.

The *x* axis is the thermoelectric cooler current normalised to I_{max}. The relevance of I_{max} is that above this current the cooling effect actually falls off. Peltier cooling increases linearly with current but unwanted ohmic heating increases as the square of the current; at I_{max}, the ohmic heating becomes greater than the Peltier cooling.

In Fig. 4, the curves of the temperature/current characteristic for heat loads are normalised to Q_{max}. The outermost curve is with no heat load so at I_{max} the temperature drop will be ΔT_{max}. This is the top right point on the graph.

The curves are then shown in steps of 0.2 for Q/Q_{max}. A finer spacing would have been better but Lotus 1-2-3 only plots six curves at a time. The situation with the heat input $Q=Q_{max}$ is represented at the bottom right corner where at I_{max} the cooling effect has been reduced to zero by the heat load.

For copyright reasons, these curves are not prepared from any one manufacturer's data. Rather they are plotted from a simple formula developed by looking at the performance of several single stage thermoelectric coolers from various manufacturers;

$$\frac{\Delta T}{\Delta T_{max}} = 2\frac{I}{I_{max}} - \left(\frac{I}{I_{max}}\right)^2 - \frac{Q}{Q_{max}}$$

As far as I know, this equation is not a recognised one, but it does have some basis in reality. The first two terms on the right hand side represent Peltier cooling and ohmic heating respectively. When differentiated with respect to current, these reduce to zero at I_{max}.

The heat loading term goes inversely as the temperature drop as would be expected from the conduction and convection terms. Errors due to other terms are ignored.

The optimum thermoelectric cooler performance is obtained on the line marked 'optimum'; this line has a 1:1 slope for normalised temperature *versus* current.

For the example used here, T/T_{max} is 0.83 so I/I_{max} becomes 0.83. The ratio Q/Q_{max} is found by interpolating between the curves and is seen to be 0.14. The required value of Q_{max} is now available as 0.905W heat input divided by 0.14, which is 6.4W.

Now you have enough information to refer to the manufacturer's literature and select a thermoelectric cooler having a ΔT_{max} of 60ºC or more and a Q_{max} of 6.4W, or slightly greater.

For example, the Marlow Industries *MI 1061* has the following characteristics,

ΔT_{max}	64ºC in dry nitrogen
Q_{max}	6.4W
I_{max}	5.3A at approximately 1.9V
Cold plate	13×15mm^2

When these numbers are normalised and plotted on the curves, the following are derived,

$\Delta T/\Delta T_{max}$	0.78
Q/Q_{max}	0.14
I/I_{max}	0.72

This is a point just above the optimum line giving the 50ºC drop for a thermoelectric cooler current of 3.8A. As a matter of interest the Marlow data sheet predicts a current of 3.7A to the accuracy to which the charts can be read.

I should point out that this thermoelectric cooler has a cold plate smaller than the ccd quoted; this will mean there is some exposed area at the back of the ccd which will behave as an extra heat leak. It may be necessary to go back and recalculate the heat input after the thermoelectric cooler has been chosen.

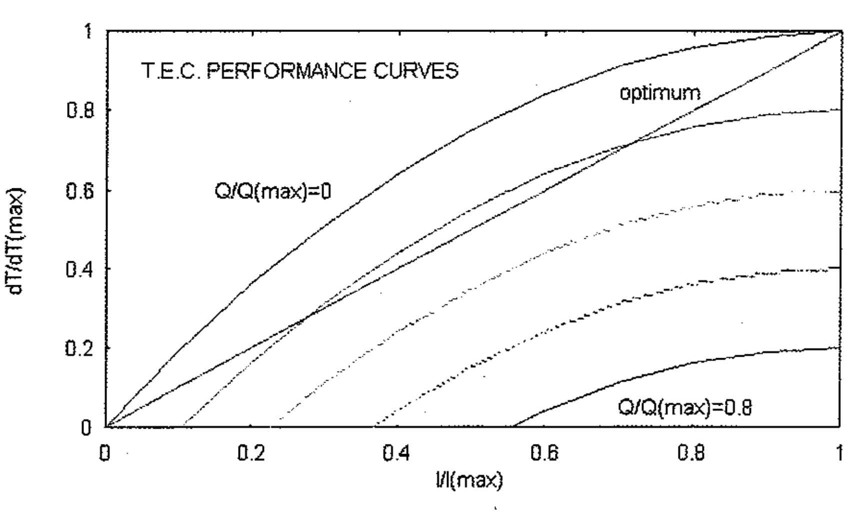

Fig. 4. Typical normalised Peltier device performance curves. Note that at I_{max}, ohmic heating becomes greater than Peltier cooling.

If there is any doubt as to whether a thermoelectric cooler will deliver the expected performance then it is essential to consult the manufacturer's data. This is especially true of the multistage units.

If you are working to a tight budget that allows you to buy only one thermoelectric cooler, it is worth getting one that can pump more heat into than is really needed in case your heat input estimate on the cold plate is too optimistic.

The *MI 1023* has Q_{max} of 9.2W, and would allow for a margin of safety.

Choosing a heat sink

Possibly the most common reason for disappointing performance with a thermoelectric cooler is inadequate provision for removing heat from the hot plate.

Manufacturers usually quote the thermoelectric cooler parameters with a hot plate temperature of 25 to 30º C and anomalies can be expected if this temperature is wildly different.

A more insidious effect is the possibility of thermal runaway if the heat sink is too small. The thermoelectric cooler becomes less efficient as the temperature differential increases; this leads to a warming of the heat sink making matters worse.

There comes a point where the hot plate warms up at a faster rate than the cold plate is cooled; increasing the thermoelectric cooler current now causes a warming of the cold plate. This results in a phase reversal in the control loop and the system locks up at full current. In this state the only option is to switch off and wait for things to cool down.

This effect is separate from the effect of running the thermoelectric cooler above I_{max} and can occur long before this current if the heat sink specification is seriously deficient.

Heat emerging from the hot plate is the sum of the dc power in and the heat entering the cold plate. To a first approximation, the thermoelectric cooler behaves as a resistor given by the volts at I_{max} which for the *MI 1061* will be 0.36Ω.

This value is not perfectly constant for all conditions; the resistance appears somewhat higher for low values of Q/Q_{max}. But the difference can be ignored if an oversized heatsink is fitted. Thus the expected heat output in the example is,

$$0.905W+3.8A^2\times0.36\Omega=6.1W$$

Knowing this number, the remainder of the problem is solved using normal heat sink design procedures. From the original problem details, the air temperature is 20ºC and the heat sink 30ºC so the maximum thermal resistance that can be accepted is 30–20ºC, i.e. 10ºC, divided by 6.1W. A heat sink capable of dissipating 1.5ºC/W or more would be a good choice.

What about water?

this is going to be quite a large heat sink, it may be worth considering a water loop. If the loop is using tap water at 15ºC and you allow 2ºC due to the thermal resistance from the hot plate to the water, the thermoelectric cooler now has to produce a ΔT of 37ºC.

Going back to the curves, you will find that the thermoelectric cooler current falls to 3A and the total heat to be disposed of is 4.2W. Remembering the heat capacity of water is 4.2 joules/cc/ºC this is a flow rate of only 1cc per second if the water is assumed to heat up by 1ºC in crossing the hot plate.

Electronic considerations

Thermoelectric coolers of the same outside dimensions and thermal properties can be made from many small p-n junctions in series, or just a few large ones. Devices in the latter category are cheaper to make and, in my experience, slightly more robust.

However, the resulting resistance can be ry low – less than an ohm as in this exam-e. This can make designing the driving electronics less straightforward – especially if a power-efficient design is required. Simple linear circuits tend to dissipate a lot of power. Switch-mode designs naturally give much better results. Remember that with very low resistance units the effects of long cable runs can seriously increase the overall power consumption.

The *MI 1013* happens to be especially convenient, with its specification of 8.5V at an I_{max} of 1A, but is quite expensive at around £60 in one-off quantities. It is also a bit small for this example. It is a single stage device with a Q_{max} of 4.8W and ΔT_{max} of 61º C.

If the thermoelectric cooler is used in a situation, where it can be required to heat or cool, then bear in mind that the devices are actually much more effective at heating than cooling.

This has implications for the gain settings in the control electronics. The servo gain can be expected to be much higher with the device heating – so beware of servo oscillations.

Mounting a Peltier device

Thermoelectric coolers are brittle, fragile and expensive. The heat sink should be milled flat and the thermoelectric cooler held down with the minimum of compression.

Even though they are reversible, devices are always mounted with the wires on the hot plate to prevent a significant heat leak. They should never be subjected to tension or shear forces, or exposed to temperatures much above 100ºC since the p-n junctions are soldered together with special low melting point solder.

If heat is removed using a water loop then some consideration should be given to protecting the thermoelectric cooler in the event of a leak or pump failure. Farnell and RS stock some useful bimetallic cutouts which are suitable if size is not a problem.

Some devices have both plates metallised with copper enabling the unit to be soldered down. This enables a very good thermal contact to be made and has the extra advantage that no supporting structure is required thus minimising heat leaks to the cold plate. Needless to say use of the manufacturer's special solder is mandatory. Heat conductive glue (RS 850-984) is an option.

Understandably, you may be reluctant to solder together an expensive thermoelectric cooler and ccd until the system is proven. It is usually possible to devise some means of compressing the ccd – or whatever device – to thermoelectric cooler and its heat sink. Inevitably though, the device used to compress the components adds to the heat input to the cold face.

If the system can be made to work in this state then results are always slightly better when the supports are removed and the components glued or soldered together. ■