CIRCUIT CELLAR®
THE MAGAZINE FOR COMPUTER APPLICATIONS

CONSIDERING THE DETAILS

Bob Perrin

The Basics of Thermocouples

This month, Bob sets out to shed some light on the mysteries of the thermocouple by explaining hot junctions, cold junctions, and dissimilar metals. As usual, he provides enough circuits and information to get you ready to design a thermocouple into your next project.

A mathematician, a physicist, and an engineer were at lunch. The bartender asked the three gentlemen, "what is this pi I hear so much about?"

The mathematician replied, "pi is the ratio of a circle's circumference to its diameter."

The physicist answered, "pi is 3.14159265359."

The engineer looked up, flatly stated, "Oh, pi's about three," then promptly went back to doodling on the back of his napkin.

The point is not that engineers are sloppy, careless, or socially inept. The point is that we are eminently practical. We are solvers of problems in a non-ideal world. This means we must be able to apply concepts to real problems and know when certain effects are negligible in our application.

For example, when designing first- or second-order filters, 3 is often a close enough approximation for pi, given the tolerance and temperature dependence of affordable components.

But, before we can run off and make gross approximations, we must understand the physical principles involved in the system we're designing. One topic that seems to suffer from gross approximations without a firm understanding of the issues involved is temperature measurement with thermocouples.

Thermocouples are simple temperature sensors consisting of two wires made from dissimilar alloys. These devices are simple in construction and easy to use. But, like any electronic component, they require a certain amount of explanation. The intent of this paper is to present and explain how to use thermocouples and how to design thermocouple interfaces.

A TAIL OF TWO METALS

Figure 1a shows a thermocouple. One junction is designated the hot junction. The other junction is designated as the cold or reference junction. The current developed in the loop is proportional to the difference in temperature between the hot and cold junctions. Thermocouples measure differences in temperature, not absolute temperature.

To understand why a current is formed, we must revert to physics. Unfortunately, I'm not a physicist, so this explanation may bend a concept or two, but I'll proceed nonetheless.

Consider a homogenous metallic

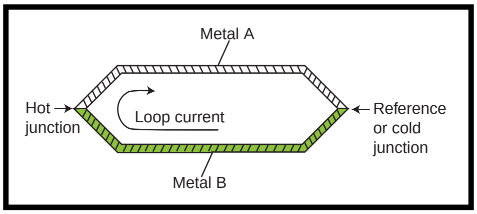

Figure 1a—*Two wires are all that are required to form a thermocouple.*

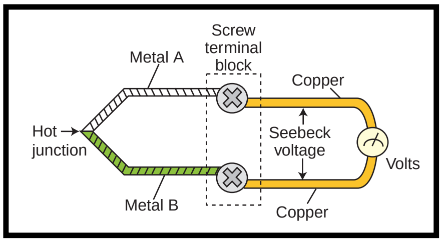

Figure 1b—*To use a thermocouple, you must have a measurement system.*

wire. If heat is applied at one end, the electrons at that end become more energetic. They absorb energy and move out of their normal energy states and into higher ones. Some will be liberated from their atoms entirely. These newly freed highly energetic electrons move toward the cool end of the wire. As these electrons speed down the wire, they transfer their energy to other atoms. This is how energy (heat) is transferred from the hot end to the cool end of the wire.

As these electrons build up at the cool end of the wire, they experience an electrostatic repulsion. The not-so-energetic electrons at the cool end move toward the hot end of the wire, which is how charge neutrality is maintained in the conductor.

The electrons moving from the cold end toward the hot end move slower than the energetic electrons moving from the hot end toward the cool end. But, on a macroscopic level, a charge balance is maintained.

When two dissimilar metals are used to form a thermocouple loop, as in Figure 1a, the difference in the two metals affinity for electrons enables a current to develop when a temperature differential is set up between the two junctions.

As electrons move from the cold junction to the hot junction, these not-so-energetic electrons are able to move easier in one metal than the other. The electrons that are moving from the hot end to the cold end have already absorbed a lot of energy, and are free to move almost equally well in both wires. This is why an electric current is developed in the loop.

I may have missed some finer points of the physics, but I think I hit the highlights. If anyone can offer a more in-depth or detailed explanation, please e-mail me. One of the best things about writing for a technical audience is learning from my readers.

BREAKING THE LOOP

If you use thermocouples, you must insert a measurement device in the loop to acquire information about the temperature difference between the hot and cold junctions. Figure 1b shows a typical setup. The thermocouple wires are brought to a terminal block and an electric circuit measures the open circuit voltage.

When the thermocouple wires are connected to the terminal block, an additional pair of thermocouples is formed (one at each screw terminal). This is true if the screw-terminals are a different alloy from the thermocouple wires. Figure 1c shows an alternate representation of Figure 1b. Junction 2 and junction 3 are undesired artifacts of the connection to the measurement circuitry. These two junctions are commonly called parasitic thermocouples.

In a physical circuit, parasitic thermocouples are formed at every solder joint, connector, and even every internal IC bond wire. If it weren't for something called the Law of Intermediate Metals, these parasitic junctions would cause us endless trouble.

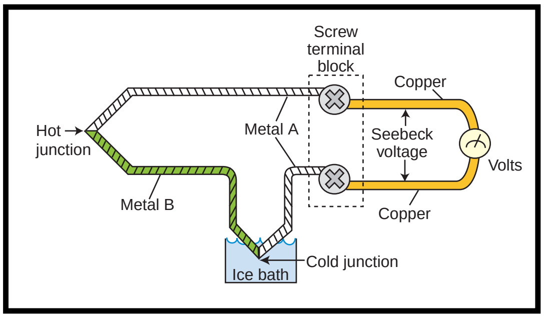

Figure 2—*By inserting a short pigtail of Metal A onto the terminal block where Metal B would normally connect, we move the cold junction.*

The Law of Intermediate Metals states that a third metal may be inserted into a thermocouple system without affecting the system if, and only if, the junctions with the third metal are kept isothermal (at the same temperature).

In Figure 1c, if junction 2 and junction 3 are at the same temperature, they will have no effect on the current in the loop. The voltage seen by the voltmeter in Figure 1b will be proportional to the difference in temperature between Junction 1 and Junctions 2 and 3.

Junction 1 is the hot junction. The isothermal terminal block is effectively removed electrically from the circuit, so the temperature of the cold junction is the temperature of the terminal block.

MEASURING TEMPERATURE

Thermocouples produce a voltage (or loop current) that is proportional to the difference in temperature between the hot junction and the reference junction. If you want to know the absolute temperature at the hot junction, you must know the absolute temperature of the reference junction.

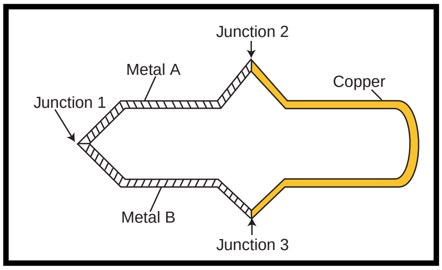

Figure 1c—*The act of connecting a measurement system made of copper introduces two parasitic thermocouples.*

There are three ways to find out the temperature of the reference junction. The simplest method is to measure the temperature at the reference junction with a thermistor or semiconductor temperature sensor such as Analog Devices' TMP03/04. Then, in software, add the measured thermocouple temperature (the difference between the hot junction and the reference junction) to the measured temperature of the reference junction. This calculation will yield the absolute temperature of the hot junction.

The second method involves holding the reference junction at a fixed and known temperature. An ice bath, or an ice slushy, is one of the most common methods used in laboratory settings. Figure 2 shows how this is accomplished.

Alternately, we could have omitted the pigtail of Metal A and just immersed the terminal block in the ice. This would work fine, but it would be

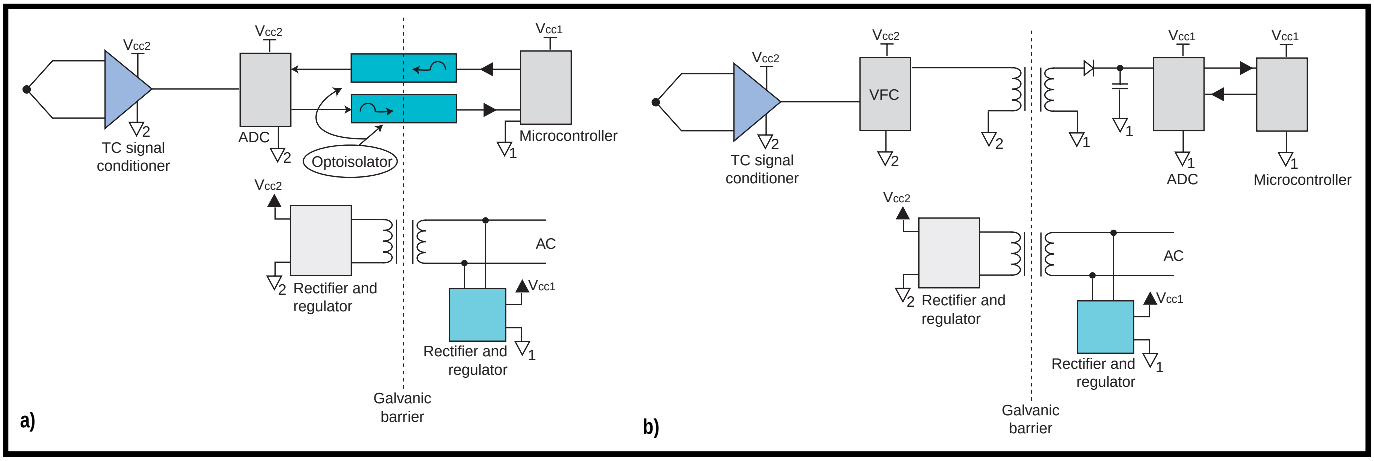

Figure 3—*Galvanic isolation to a few thousand volts is easy (but a little expensive) using opto-isolation (**a**) and inexpensive (but a bit more challenging) using a VFC and a transformer (**b**).*

much messier than the method shown in Figure 2.

Sometimes, the temperature of the cold junction (terminal block) in Figure 1c is allowed to float to ambient. Then ambient is assumed to be "about 25°C," or some other "close enough" temperature. This method is usually found in systems where knowing the temperature of the hot junction is not overly critical.

The third method used to nail down the cold junction temperature is to use a cold junction compensation IC such as the Analog Devices AD594 or Linear Technology LT1025. This method sort of combines the first two methods.

These ICs have a temperature sensor in them that detects the temperature of the cold junction. This is presumably the same temperature as the circuit board on which the IC is mounted. The IC then produces a voltage that is proportional to the voltage produced by a thermocouple with its hot junction at ambient and its cold junction at 0°C. This voltage is added to the EMF produced by the thermocouple. The net effect is the same as if the cold junction were physically held at 0°C.

The act of knowing (or approximating) the cold junction temperature and taking this information in to account in the overall measurement is referred to as cold junction compensation. The three techniques I discussed are each methods of cold junction compensation.

The ice bath is probably the most accurate method. An ice slushy can maintain a uniformity of about 0.1°C without much difficulty. I've read that an ice bath can maintain a uniformity of 0.01°C, but I've never been able to achieve that level of uniformity. Ice baths are physically awkward and therefore usually impractical for industrial measurements.

The off-the-shelf cold junction compensation ICs can be expensive and generally are only accurate to a few degrees Celsius, but many systems use these devices.

Using a thermistor, or even the PN junction on a diode or BJT, to measure the cold junction temperature can be fairly inexpensive and quite accurate. The most common difficulty encountered with this system is calibration. Prudent positioning of the sensor near, or on the terminal block is important.

If the terminal block is to be used as the cold junction (see Figure 1b), the terminal block must be kept isothermal. In practice, keeping the terminal block truly isothermal is almost impossible. So, compromises must be made. This is the stock and trade of engineers. Knowing what is isothermal "enough" for your application is the trick.

Lots of money can be wasted on precision electronics if the terminal block's screw terminals are allowed to

Type	Materials: Positive side	Materials: Negative side	Usable temperature range in degrees Celsius
B	Pt + 30% Rh	Pt + 6% Rh	0 to 1820
E	Ni + 10% Cr Chromel*	Cu + 43% Ni Constantan	-270 to 1000
J	Fe Iron	Cu + 43% Ni Constantan	-210 to 1200
K	Ni + 10% Cr Chromel*	Ni + 2% Al + 2% Mn + 1% Si Alumel*	-270 to 1372
N	Ni + 14% Cr + 1.5% Si Nicrosil	Ni + 4.5% Si + 0.1% Mg Nisil	-270 to 1300
R	Pt + 13% Rh	Pt	-50 to 1768
S	Pt + 10% Rh	Pt	-50 to 1768
T	Cu Copper	Cu + 43% Ni Constantan	-270 to 400

*Chromel and Alumel are trademarks of Hoskins Manufacturing Company

Table 1—*There are a wide variety of industry-standard alloy combinations that form standard thermocouples. The most commonly used are J, K, T, and E.*

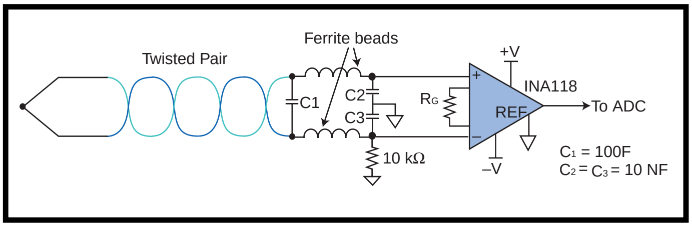

Figure 4—*The common-mode filter and common-mode rejection characteristics pay off in thermocouple amplifiers.*

develop a significant thermal gradient. This condition generally happens when power components are placed near the terminal blocks. You must pay careful attention to keeping the temperature stable around the terminal blocks.

There are two broad classes of temperature-measurement applications. The first class involves measuring absolute temperature. For example, you may want to know the temperature of the inside of an oven relative to a standard temperature scale (like the Celsius scale). This type of application requires that you know precisely the absolute temperature of the reference junction.

The second type of measurement involves measuring differences in temperature. For example, in a microcalorimeter, you may want to measure the temperature of the system, then start some chemical reaction and measure the temperature as the reaction proceeds. The information of value is the difference between the first measurement and the subsequent ones.

Systems that measure temperature differences are generally easier to construct because control or precise measurement of the reference junction isn't required. What is required is that the reference junction remain at a constant temperature while the two measurements occur. Whether the reference junction is at 25.0°C or 30.0°C isn't relevant because the subtraction of consecutive measurements will remove the reference junction temperature from the computed answer.

You can use thermocouples to make precise differential temperature measurements, but you must ensure the terminal block forming the cold junction is "close enough" to isothermal. You must also ensure that the cold junction has enough thermal mass so it will not change temperature over the time you have between measurements.

PRACTICAL MATTERS

Thermocouples are given a letter designation that indicates the materials they are fabricated from. This letter designation is called the thermocouples "type." Table 1 shows the common thermocouples available and their usable temperature ranges.

Each thermocouple type will produce a different open-circuit voltage (Seebeck voltage) for a given set of temperature conditions. None of these devices are linear over a full range of temperatures. There are standard tables available that tabulate Seebeck voltages as a function of temperature [1]. There are also standard polynomial models available for thermocouples.

Thermocouples produce a small Seebeck voltage. For example, a type K thermocouple produces about 40 µV per degree Celsius when both junctions are near room temperature. The most sensitive of the thermocouples, type E, produces about 60 µV per degree Celsius when both junctions are near room temperature.

In many applications, the range of temperatures being measured is sufficiently small that the Seebeck voltage is assumed to be linear over the range of interest. This eliminates the need for lookup tables or polynomial computation in the system. Often the loss of absolute accuracy is negligible, but this tradeoff is one the design engineer must weigh carefully.

CIRCUITS

When designing a thermocouple interface, there are only a few pieces of information you need to know:

- what type of thermocouple will be used
- what is the full range of temperatures the hot junction will be exposed to
- what is the full range of temperatures the cold junction will be exposed to
- what is the temperature resolution required for your application
- does your system require galvanic isolation
- what type of cold junction compensation will be used

If the answer to the last question requires the analog addition of a voltage from a commercial cold junction compensation IC, then the manufacturer of the IC will probably supply you with an adequate reference design. If you plan to do the cold-junction compensation either physically (by an ice bath) or in software (by measuring the cold junction's temperature with another device), then you must build or buy a data acquisition system.

Galvanic isolation is an important feature in many industrial applications. Because thermocouples are really just long loops of wire, they will often pick up high levels of common-mode noise. In some applications, the thermocouples may be bonded to equipment that is at line voltage (or higher).

In this case, galvanic isolation is required to keep high-voltage AC out of your data acquisition system. This type of isolation is usually accomplished in one of two ways—using either an opto-isolator or a transformer. Both systems require the thermocouple signal conditioner to allow its ground to float with respect to earth ground. Figure 3a and 3b outlines these schemes.

Because the focus of this article is on the interface to the thermocouple, I'll have to leave the details of implementing galvanic isolation to another

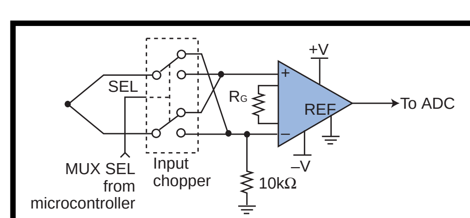

Figure 5—*An input chopper like a CD4052 is all that is necessary to null signal conditioner offsets.*

article.

Given the tiny voltage levels produced by a thermocouple, the designer of the signal-conditioning module should focus carefully on noise rejection. Using the common-mode rejection (CMR) characteristics of a differential amplifier is a good place to start. Figure 4 shows a simple yet effective thermocouple interface.

The monolithic instrumentation amplifier (in-amp) is a \$2–\$5 part (depending on grade and manufacturer). These are usually 8-pin DIP or SOIC devices. In-amps are simple differential amplifiers. The gain is set with a single external resistor. The input impedance of an in-amp is typically 10 Gigaohms.

Certainly you can use op-amps, or even discrete parts to build a signal conditioner. However, all the active components on a monolithic in-amp are on the same dice and are kept more-or-less isothermal. This means in-amp characteristics behave nicely over temperature. Good CMR, controllable gain, small size, and high input impedance make in-amps perfect as the heart of a thermocouple conditioning circuit.

Temperature tends to change relatively slowly. So, if you find your system has noise, you can usually install supplementary low-pass filters. These can be implemented in hardware or software. In many systems, it's not uncommon to take 128 measurements over 1 s and then average the results. Digital filters are big cost reducers in production systems.

Another problem often faced when designing thermocouple circuits is nulling amplifier offset. You can null the amplifier offset a variety of ways [2], but my favorite is by chopping the input. Figure 5 shows how this process can be accomplished.

Thermocouples have such small signal levels, gains on the order of 1000 V/V are not uncommon, which means an op-amp or in-amp with a voltage offset of even 1 mV will have an offset at the output on the order of volts.

The chopper in Figure 5 allows the microcontroller to reverse the polarity of the thermocouple. To null the circuit, the microcontroller will take two measurements then subtract them.

First, set the chopper so the ADC measures GAIN (Vsensor + Voffset). Second, set the chopper so the ADC measures GAIN (–Vsensor + Voffset).

Subtract the second measurement from the first and divide by two. The result is GAIN*Vsensor. As you can see, this is exactly the quantity we are interested in. The in-amp's offset has been removed from the measurement.

CLOSING TIME

In 1821, Thomas J. Seebeck discovered that if a junction of two dissimilar metals is heated, a voltage is produced. This voltage has since been dubbed the Seebeck voltage.

Thermocouples are found in everything from industrial furnaces to medical devices. At first glance, thermocouples may seem fraught with mystery. They are not. After all, how can a device that's built from two wires and has been around for 180 years be all that tough to figure out?

When designing with thermocouples, just keep these four concepts in mind and the project will go much smoother. First, thermocouples produce a voltage that is proportional to the difference in temperature between the hot junction and the reference junction.

Second, because thermocouples measure relative temperature differences, cold junction compensation is required if the system is to report absolute temperatures. Cold-junction compensation simply means knowing the absolute temperature of the cold junction and adjusting the reported temperature value accordingly.

The third thing to remember is that thermocouples have a small Seebeck voltage coefficient, typically on the order of tens of microvolts per degree Celsius. And last, thermocouples are non-linear across their temperature range. Linearization, if needed, is best done in software.

Armed with these concepts, the circuits in this article, and a bit of time, you should have a good start on being able to design a thermocouple into your next project.

Over the last ten years, Bob has designed instrumentation for agronomy, soil physics, and water activity research. He has also designed embedded controllers for a variety of other applications. For more technical resources and articles, visit Bob's online library at www.engineerbob.com. You may reach Bob with comments and questions at bob@engineerbob.com.

REFERENCES

[1] www.omega.com/pdf/temperature/Z/zsection.asp

[2] B. Perrin, "Practical Analog Design," *Circuit Cellar*, 94, May 1999.

SOURCES

AD594, **TMP03/04**
Analog Devices
(617) 329-4700
Fax: (617) 329-1241
www.analog.com

INA118
Burr-Brown Corp.
(520) 746-1111
Fax: (520) 889-1510
www.burr-brown.com

LT1025
Linear Technology
(408) 432-1900
Fax: (408) 434-0507
www.linear-tech.com