Tellurex

1462 International Drive
Traverse City, MI 49686
231-947-0110
www.tellurex.com

An Introduction To Temperature Control Of Thermoelectric Systems

Introduction:

One of the attractive features of thermoelectric (TE) technology, is that it offers an incredible degree of controllability. With a properly 'tuned' controller, it is possible to maintain systems well within 0.1° C of set point. Unfortunately, many off-the-shelf solutions for temperature control are not well suited to the thermoelectric world because they were designed for heating or cooling hardware that is very different—and often far less responsive—than TE devices. This leaves designers groping for alternatives. This guide is intended to offer designers some practical guidance in exploring the vast range of possibilities.

Types of Control

Basically there are two types of temperature control: *thermostatic* and *steady-state*.

Thermostatic Control

With thermostatic control, a thermal load is maintained between two temperature limits. For example, in a cooling application, the controller may energize TE cooling power when the thermal load rises to 30° C, then turn off cooling power when the temperature cools to 27° C; the system would, therefore, continually vary between 27 and 30° C (see *Figure 1*). The difference between the high and low temperature limits, would be the system's 'hysteresis'—in this case, 3° C. Thermostatic control is often the least costly alternative and is worth considering whenever a user can tolerate some appreciable variation in operating temperatures.

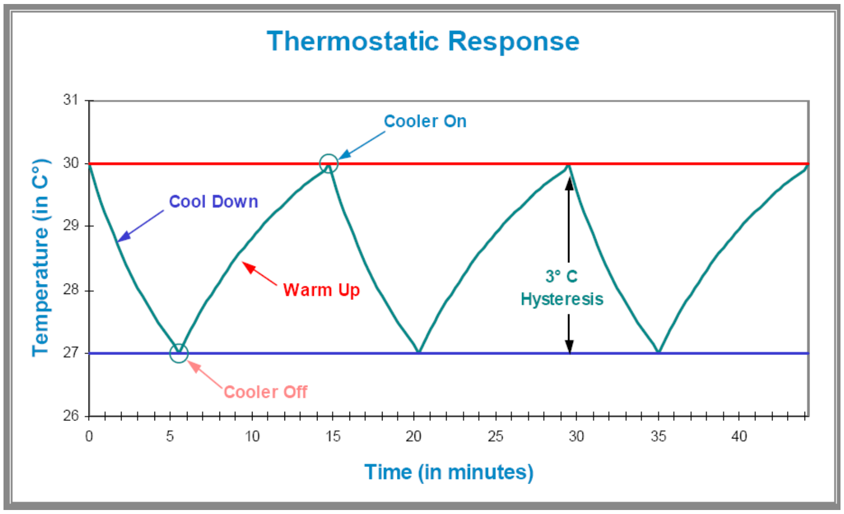

Figure 1

In the world of thermoelectric technology, thermostatic control is usually accomplished by switching TE power *electronically* rather than with electro-mechanical devices such as relays. This is due to two principal reasons.

1. The devices are switching *direct* rather than alternating current; this makes mechanical contacts more vulnerable to the pitting and premature wear which result from arcing.

2. Using electronic means, you don't have to worry about ratings for the typical number of switching operations before failure—you can switch electronically as often as you want.

Because of this, it is easier to provide thermostatic control with minimal hysteresis. In some systems, the hysteresis can be targeted as low as 1° C.

WARNING: It is *not* recommended that bimetal devices such as snap disks be employed for direct electro-mechanical switching of thermoelectric modules. While these components can be used to energize relays for thermostatic control, the *direct* switching of thermoelectric current *frequently* leads to permanent closure of bimetal devices or shifting of set points. This can potentially lead to *serious* safety issues and TE device destruction in heating applications. Do not rely upon the assurances of a bimetal device manufacturer on reliability of his/her product in switching high-current DC. Such use with TE's is risky in *all* cases!

The component of choice for *electro*-thermostatic switching with thermoelectric technology, is the power MOSFET (i.e., metal oxide semiconductor field effect transistor). Relatively inexpensive power MOSFET's can be found with 'on' resistances as low as 0.01 ohms. This means that very little power is lost in the switching device. For example, if you were driving two typical TE devices connected electrically in parallel and collectively drawing 10 amps, there would only be 1 Watt ($P=I^2R$) dissipated in this power MOSFET with a voltage drop of just 0.1 VDC ($V=IR$). Of course, some designers naturally gravitate toward solid state relays in this sort of situation, but these devices are seldom used with thermoelectric technology. SSR's are much more resistive than MOSFET's, drop more voltage, dissipate a lot more power, and are bulkier. The only reason to use a solid state relay with TE's, is because you don't know better.

There are times when thermostatic control of TE systems is inadvisable. For example, if the required hysteresis is too small, the system may have to cycle on and off in a matter of seconds to maintain the system within the desired range. While the controller can certainly meet these requirements, this sort of cycling is mechanically stressful to the TE modules and will lead to their premature aging and failure. This is due to differential rates of expansion and contraction among the various materials used to construct the module. Thermostatic control of TE systems should ideally be limited to applications where the controller will cycle over a period of at least a few minutes or more.

Steady-State Control

Whenever a system must be maintained within fairly tight limits, some form of steady state controller should be considered. A steady-state controller is designed to continually hold a thermal load at the set-point temperature with very little variation around the set-point. If the steady-state condition is disrupted by some sudden change in ambient conditions, the controller will quickly bring the thermal load back to a steady-state condition (provided that the system has sufficient heating or cooling capacity). To achieve steady-state control with thermoelectric technology, the controller must vary the amount of voltage or current to the TE modules so that there is just enough power to maintain the system at the desired temperature. To accomplish this, the controller must be able to make instantaneous adjustments in TE power in response to changes in the ambient environment or thermal load.

It should be noted that many system designers make the mistake of thinking that they can achieve steady-state temperature control by simply adjusting the TE power supply to a level which yields the desired temperature. This approach can only work if ambient conditions are absolutely constant. Without real temperature control, any variations in the operating environment will cause changes in the temperature of the load. The key to effective steady-state control, is that the amount of TE power be made dependent upon the ambient and thermal load conditions which exist at any particular moment.

Steady-state control is usually achieved with some variant of a proportional controller. Proportional controller design originally grew from the reality that set-point potentiometers and temperature sensors used electrical voltages as indicators of temperature. In its most simple form, a proportional controller merely amplified the difference between the set-point and sensor voltages and provided an output (in this case, TE power) which was proportional to that difference. Thus the amount of TE power was a function of the difference in temperature between the desired set-point and the measured temperature of the thermal load. For example, the set-point voltage in a particular system might be 2.73 V; if the temperature sensor developed a voltage of 2.83 V and the proportional amplifier had a gain of 100, the proportional amplifier would have an output of 10 VDC:

$$(2.83\ \text{V} - 2.73\ \text{V}) * 100 = 10\ \text{VDC}$$

The controller would then translate this amplified error voltage (i.e., 10 V) into a proportional amount of TE power. In a simple proportional controller, therefore, the precise amount of TE power at any given point in time, is a function of the amplified error between the set-point and sensor voltages. If a source of heat in the ambient environment causes greater error, the amount of TE cooling power immediately increases, then gradually decreases as the system dips toward the set-point again.

Recognize that, with a basic proportional controller, *some error must exist* between the set-point and sensor voltages in order to have any power provided to the TE modules—if there is no error, there is no voltage difference to amplify, and as a result, no TE power generated by the control system. In practice, a simple proportional controller will bring the system to a temperature where there is just enough error—just enough amplified error voltage—to yield the level of TE power required to stay at that temperature. For example, in a given ambient environment, if:

1) 2 V of amplified error voltage is required for the controller to provide enough cooling power to maintain a 15° C temperature,

2) it takes a 0.1° C temperature difference to create an amplified error voltage of 2 V, and

3) the set-point is adjusted to a voltage corresponding to 15° C,

The system will settle to a steady-state temperature of 15.1° C. If the ambient temperature rises, it will take more cooling power—and thus more error—to maintain the temperature near the set-point, and the system will settle to a somewhat higher temperature. If the ambient temperature decreases (but remains above the set-point), it will take less cooling power (and less error) to maintain the temperature near the set-point—the system will then settle to a temperature which is closer to the set point. The only way that there will be no error between the set-point and sensor temperatures, is if no TE power is required to reach and maintain the set-point.

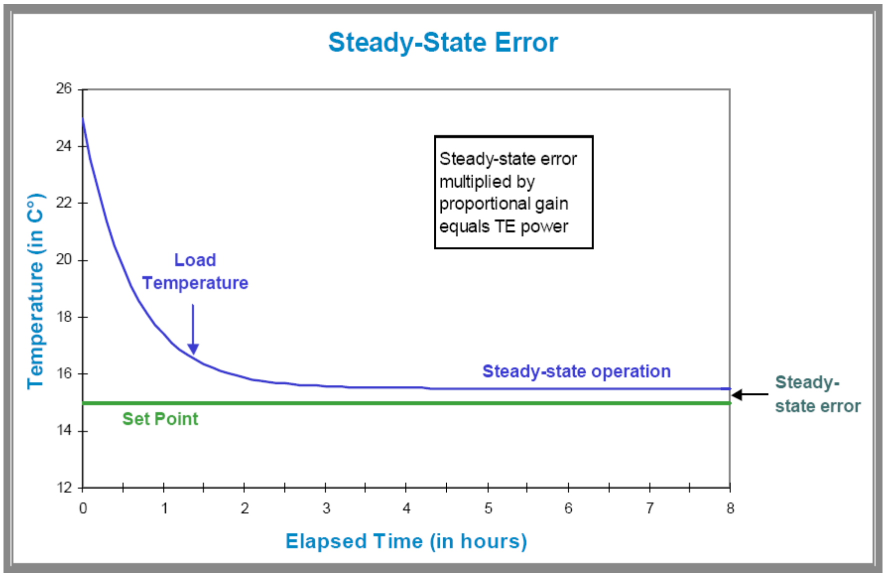

Figure 2

The amount of error which exists when the system reaches a steady state, is called 'steady-state error' (see *Figure 2*). It should be noted that the greater the gain of the proportional amplifier, the less steady-state error will be required to generate a given level of TE power; we can thus decrease steady-state error by maximizing proportional gain. It might appear, therefore, that a designer would want infinite proportional gain to decrease steady-state error to zero. Excessive proportional gain, however, makes a temperature control system too sensitive and causes it to over-respond to changes in conditions (see *Figure 3*). This results in system oscillation around the set-point, with the temperature swinging back and forth continually. Generally, the proportional gain will be set so that steady-state error is minimized while assuring stable operation of the system.

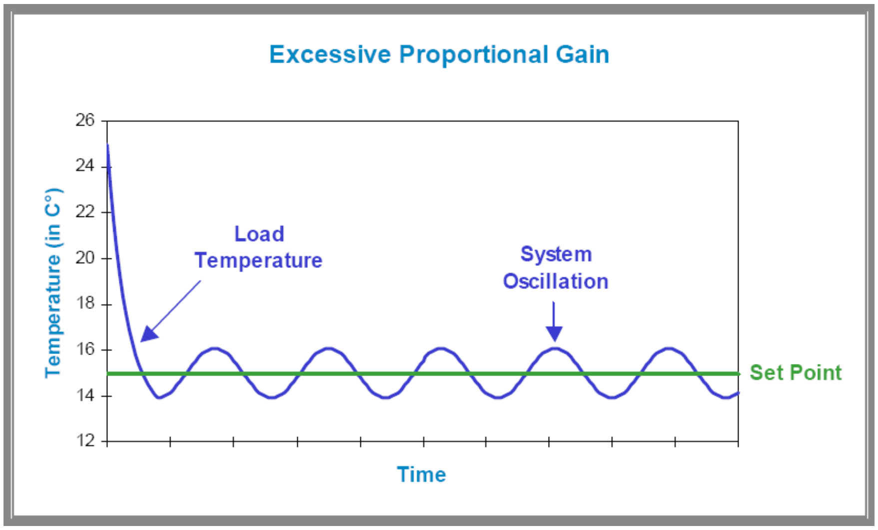

Figure 3

Note that proportional control can only approximate steady-state operation. In reality (assuming that sufficient power is available to keep the thermal load near the setpoint), the temperature of the system will fluctuate a bit as ambient conditions change (see *Figure 4*). It not only takes time to respond to any disruptions in the operating environment, but it requires different amounts of error voltage to bring the system back into balance. Assuming that the system always has enough capacity to reach a near steady-state condition, the system will vary anywhere between the set-point (where there is no TE power) and the temperature at which the steady-state error yields full TE-power. In many TE systems using proportional control alone, maximum steady-state error will only be a few tenths of a degree and this extent of inaccuracy can be easily tolerated. This is often the case, for example, when controlling the air temperature of an enclosed space. In other systems, however, the maximum steady-state error may be many degrees—a level of potential error which is seldom acceptable. This might be the case, for example, when the system is directly controlling the temperature of liquids or metallic solids. In these cases, because temperature responds so quickly to changes in TE current, the proportional gain may have to be decreased to a point where steady-state error becomes substantial. In situations where the level of steady-state error is clearly unacceptable, additional control functions must be added to compensate.

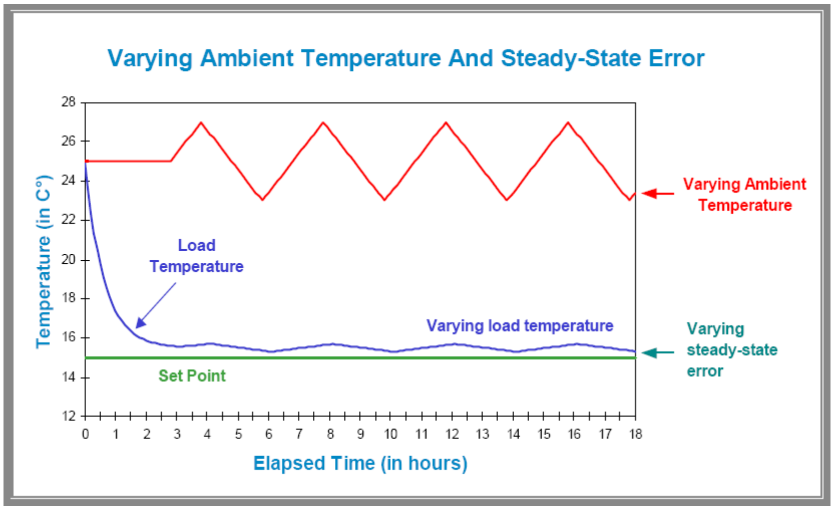

Figure 4: With a fluctuating ambient, while the proportional gain minimizes the amount of temperature change in the load, some variation still occurs.

It is with the addition of an *integral* amplifier that we can virtually eliminate steady-state error in a proportional control system. The integral amplifier essentially integrates the amplified error voltage; this means that its output keeps increasing as long as an error voltage is present. The amount of increase is a function of the extent of error and the gain of the integral amplifier. Once the amplified error is eliminated, the integral amplifier simply holds its output steady until a new error voltage becomes present. The voltage output of the integral amplifier is added to the output of the amplified error voltage from the proportional amplifier to determine the amount of TE power required. (Note that if the system overshoots the set point, the integral amp can integrate in the opposite polarity to provide corrective action).

In theory, upon system power up, the amplified error voltage from the proportional amplifier will contribute to the bulk of TE power while the integral amplifier starts at zero and begins ramping up. Ultimately, the output of the integral amp exceeds than that from the proportional side, and becomes the driving force in sustaining the necessary level of TE power. Because it keys on error voltage, the integral function will ultimately ramp to a value that results in zero error. (In reality, some amount of error will still be evident because of offset errors in amplifiers and other subtle error-producing phenomena in the circuitry, but these discrepancies will be inconsequential in most systems). The higher the gain of the integral amplifier, the faster it can respond to error potential. If the integral gain is inadequate, it will be very slow to compensate; if the integral gain is too high, it will tend to over-correct and the system will oscillate. Typically the gain of the integral amplifier is set so that upon power-up, the system will 'ring' one time (i.e., it will overshoot the set-point, then undershoot it, then settle to the setpoint).

A related problem in many TE systems, is that it takes a prolonged period of time to reach the set-point. In these cases, an integral amplifier can 'saturate' and 'lock-up' at full output. This can cause the TE system to overshoot the set-point by a considerable amount and then stay in this condition for several minutes until the amplifier gets out of its locked-up state;

only then can it integrate again to bring the system back toward the set-point. For this reason, it is usually desirable for the integral amplifier to have "anti-windup" circuitry; this inhibits integral action until the system gets close to the set-point temperature. Typically, this will prevent the integral amp from saturating.

A controller which uses only proportional and integral amplifiers, is known as a "PI" controller. Unfortunately, while integral amplifiers can effectively correct for steady-state errors, they tend to make the control system more unstable. For this reason, a derivative amplifier is often employed to bring things into more optimal control. Where the proportional amplifier keys only on the amount of error, a derivative amplifier responds to the *rate of change*; the faster that the temperature changes, the greater is the output of the derivative amplifier. It is thus used to provide an anticipatory response to sudden variations in load or ambient conditions.

Let's look at an example. A cold plate is being cooled and has reached a steady-state temperature of 5° C. You suddenly place your hand on the plate which causes the temperature to rise quickly. If you only have proportional control, the TE output would simply be a function of the amount of error which results. With the addition of the derivative component, however, the rapid change of temperature brings additional cooling power on line to respond to the disruption in a more forceful and effective way. Note, too, that as the system is being brought back to the set-point, if it starts cooling too quickly, the derivative amplifier will apply ‘the brakes’ and decrease cooling power, thus minimizing or preventing overshoot. The higher the gain of the derivative amplifier, the greater will be its instantaneous response to change. The gain must be set to a level which effectively minimizes the effect of system disruptions, but not so high that it tends to over-compensate or the system will oscillate. High derivative gain will also tend to slow the system down in initially ramping to the set-point; in some cases this is desirable or even necessary, in others, it is not. A controller which uses proportional, integral, and derivative functions, is known as a "PID" controller and this is the most common form of steady-state control.

It is worth mentioning here that most off-the-shelf PID controllers designed for other heating and cooling technologies, tend to provide problematic derivative response for thermoelectric systems. This is particularly true with cold plates and liquid recirculating heat exchangers where TE systems are especially responsive. Other technologies are typically far more sluggish and this is reflected by the derivative gain settings offered by off-the-shelf controllers. All too frequently, there will be minimal useful adjustment range for TE’s on derivative action.

With steady-state control, one of two approaches is taken to vary the amount of power provided to the TE device: 1) pulse-width modulation (a.k.a., ‘PWM'), and 2) linear control. Pulse-width modulation provides the most efficient means, because the driving transistor (usually a power MOSFET) is either fully on or fully off. In either state of conduction, very little power is dissipated within the transistor. With PWM, the TE modules are pulsed on and off at a specific frequency with the controller merely changing the percentage of ‘on’ time vs. ‘off’ time (i.e., the ‘duty cycle’) during the pulse train (see *Figure 5*). One of the problems in providing PWM for TE devices, is that many off-the-shelf PID controllers will provide pulses at only 1 Hz (i.e., cycle per second) or less, and this low frequency drive is more thermally stressful to the TE modules than pulsing at higher frequencies. At Tellurex, we recommend that TE modules be pulsed at 60 Hz or greater when using PWM to minimize the effects of thermal expansion and contraction. Fortunately, controllers are available which have analog outputs; with a fairly simple interface circuit (see *Figure 6*), the analog level can be converted to a proportional pulse-width modulated drive signal at the desired frequency. (Note that when using PWM, the 'on' voltage should never be allowed to go beyond the V_{MAX} specification for the TE module or performance will suffer.) A small microcontroller can also be used for this purpose.

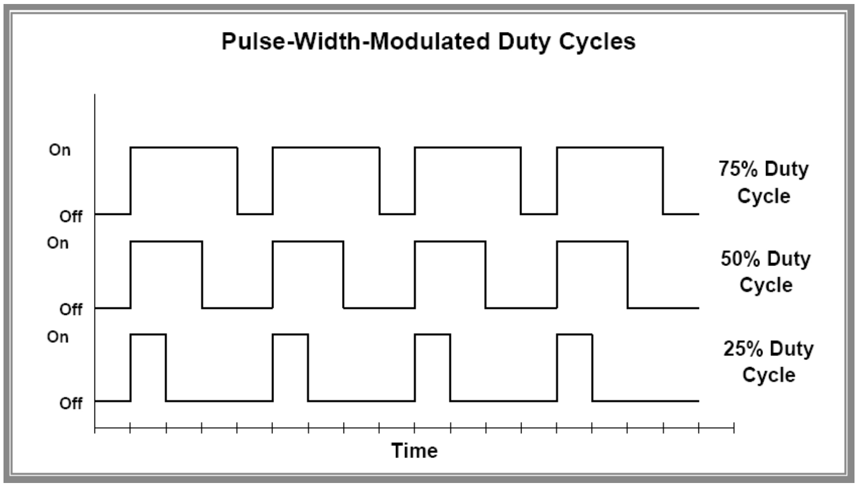

Figure 5

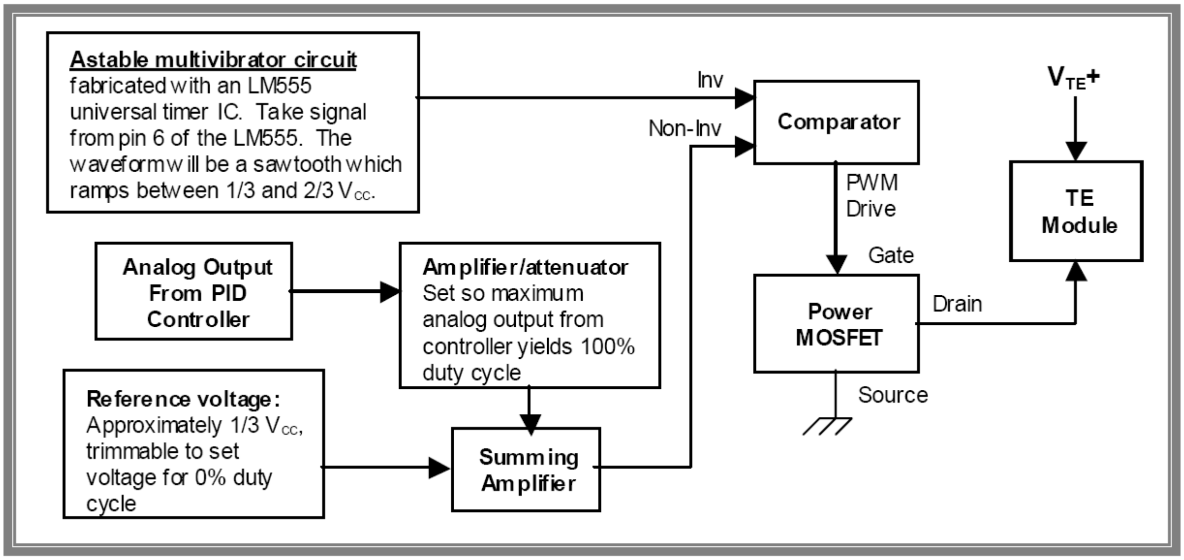

Figure 6

Despite its efficiency advantages, PWM cannot be used for steady-state control in some cases because of the potential for electromagnetic interference (EMI) from the high current TE pulses and/or because poor transient response in the available power supply can cause audible high frequency noise. In these instances, some sort of linear drive circuit is created to provide a variable DC level for powering the TE module. Most designers employ bipolar junction transistors (BJT's) in these instances to provide a linear drive current for the TE modules. Unfortunately, these BJT drivers can dissipate a lot more power than pulse-width-modulated MOSFET's. For example, if you use a TE module which is rated for 5 amps, 12 VDC at full power, the worst case for the BJT drive transistor would be when 6 V is provided to the TE module at 2.5 amps. Here the TE device would dissipate 15 W—as would the drive transistor. The BJT driver, therefore, would have to be mounted to a heat sink which can effectively manage 15 W of heat. Compare that to a power MOSFET with a 0.01 ohm 'on' resistance running at 50% duty cycle (the equivalent in PWM); here the driver would only dissipate an average of 0.125 Watts. It's also important to note in this example, that if multiple TE modules were used, there would be a worst case power dissipation in the BJT driver of 15 W for every TE module employed. Unfortunately, in some applications, linear drive is the only reasonable alternative, and the electrical inefficiencies must be tolerated in order to serve more important objectives. Designers should not be too quick to assume that they will require linear drive, however. With properly shaped waveforms, pulse width modulated designs have been used in automotive applications where they have easily passed electromagnetic coupling tests in certified labs.

In recent years, Tellurex electrical engineers have pioneered a means of creating linear drive using power MOSFETs. While this still is not as efficient as PWM, the much better saturation conductivity of the MOSFET, makes it far more suitable than the traditional BJT driver. This is seen when the TE system requires full power from the cooling module. In this circumstance, less power is dissipated in the drive transistor, and the voltage drop across the driver is significantly lower than with a BJT. When using a properly-selected MOSFET, all but a small fraction of a volt from the available supply, can be applied to the TE device for full-power operation. In a BJT-controlled circuit at saturation, a full volt may be dropped across the transistor resulting in a many-fold increase in power dissipation within the driver. Thankfully, too, only minimal additional circuitry is required for this improved solution to linear drive.

Digital Controllers

The preceding discussion has focused upon controllers built from analog electronic components; this is where the technology originated and it is much easier to conceptualize from that perspective. Most of the available control hardware, however, is now digitally based using microprocessors or micro-controllers. Digital PID controllers use software algorithms to interpret the differences between the measured temperature and the set point. It is software which amplifies error, integrates, and derives the proper amount of TE power to hold the thermal load in steady-state operation. Digital PID controllers basically mimic the functioning of their analog equivalents.

There are many advantages with digital controllers. Perhaps their greatest strength comes into play when a number of identical TE systems are to be fabricated; once the correct tuning has been determined for the prototype, the gain settings can be easily programmed into each controller to assure consistent performance among all units. In contrast, unless the gains are locked in with fixed resistor values, analog controllers typically employ trim pots for the gain settings, and every unit must be tuned individually. Another attractive feature offered by many digital controllers, is an LED or LCD display that can show the set-point and measured temperatures. A lot of off-the-shelf controllers can also accommodate a variety of sensor options, and many can even tune themselves to provide satisfactory performance. It is not uncommon, too, for these circuits to offer a choice between steady-state and thermostatic control (although steady state will usually win out in TE systems). Increasingly, designers are employing custom circuits based on ‘reduced instruction set controllers’ (RISC’s) which not only keep temperatures stable, but manage other system functions, as well.

The pros and cons over digital & linear approaches have shifted appreciably in recent years. At one time, digital was clearly the more expensive approach, but as RISC and other microprocessor prices have dropped dramatically, digital strategies are becoming more commonplace in lower cost controllers. Still, there are situations where analog approaches can yield highly stripped down, fixed-gain controllers more cheaply, especially in high-volume OEM applications. In many cases, too, well-tuned analog designs can still deliver superior performance over a digital approach because they respond to conditions instantaneously without any delays for processing, polling, or analysis. However, the ability of digital technology to handle so many other complexities (e.g., ramp & soak, displays, auto-tuning, repeatability of set point ranges, computer interfacing, tuning for mass deployment, etc.) give it many overriding advantages.

The Cost of Control

In virtually every TE application where temperature control is required, cost of control is a major consideration. Unfortunately, many TE system designers unconsciously drive up costs by convincing themselves that they require what is merely desired. The reality is that, for every 'inch' of quality, precision, or flexibility in controller performance, a price must be paid and sometimes the charge for a small feature is disproportionate. Wise designers will—very early in the process—examine their control objectives and separate those items which reflect necessity, from those which would simply be "nice to have". If minimal cost is important, designers must make every possible sacrifice in system specifications and avoid the temptation to add unnecessary performance requirements or superficial 'bells and whistles'. Among the most important considerations are these:

- *Is an adjustable set-point required*? The simplest and least costly controllers will be factory-calibrated for a single set-point with no method of re-adjustment (other than perhaps recalibration) for the end user. Providing a set-point range for end-user adjustment will add cost to the control system with the full extent of that increase dependent upon other control or user-interface objectives.

- *Is there a need for visual indicators for set-point or temperature*? The simplest and least costly controllers will have a 'blind' set-point range (i.e., there is no visual indication of the precise set-point temperature) and no display of the resulting temperature. For example, the set-point knob might be labeled with an arrow and the word "Cooler" to show the user which direction to turn the knob for adjustment; the user would then monitor the temperature with a thermometer and adjust as necessary, until the desired performance is achieved. If some visual indication of set-point is required, it can be provided in the least costly manner by using a printed adjustment scale and indicator knob. The addition of this feature, however: 1) may impact upon the choice of temperature sensor, and 2) will require additional calibration to assure modest accuracy in the correspondence between the knob setting and the actual set-point which results.

  The choice of a digital display will typically add great cost to the controller. This is due not only to the cost of display-related hardware, but also because some modest sources of error can no longer be effectively 'hidden' from the user. As a result, the controller manufacturer will need to go to extra lengths to 'trim out' error.

- *How much accuracy and repeatability is required*? One of the biggest challenges in designing electronic controllers, is in coping with the variations of components within tolerance. For example, while a resistor with a tolerance rating of 1% and a temperature coefficient of 50 ppm may be non-problematic in most electronic circuits, in an instrumentation or control application, the variations can result in set-point shifts of 5° C or more. Similarly, offset errors in operational amplifier IC's are inconsequential in many applications, but can contribute to serious inaccuracies in instrumentation and control circuits. This reality means that control designers sometimes must rely upon extreme measures to achieve accuracy and repeatability, and the cost of a controller will correlate highly with the degree of conformance demanded. The difference in cost between a controller which is accurate to within 2-3° C and one accurate to less than 0.1° C, is immense. As requirements for accuracy and repeatability tighten, the controller will need more and more provisions for calibration and error reduction. In extreme cases, some of the components on the control board may require temperature regulation themselves to minimize variations.

In coming to terms with these issues, designers should avoid falling into the trap of specifying absurd performance requirements. For example, some people will demand stability and accuracy to within 0.01° C of set-point when they: 1) can't dependably and accurately measure to that standard, and/or 2) will have gradients across the thermal load that exceed this specification. Just ask for what you really need and you will keep costs down substantially. Don't be tempted to pad your requirements with an additional margin of error—and if you deal with controller designers, ask them to do the same. If everyone substantially pads their performance standards, the resulting controller can cost a great deal more than necessary in the end.

- *Will the system be single-mode (i.e., just heating or just cooling) or dual mode (i.e., switchable between heating and cooling modes)*? Of course, one of the attractive features of TE technology, is that it can be used for either heating or cooling and it is often the technology of choice in situations where either mode of operation may be required depending upon ambient conditions. Dual mode control, as one might expect, comes at a higher price, particularly in steady-state systems or those requiring automatic switching between modes. Significant additional hardware is required in these instances, including relays or a transistor bridge for polarity switching. This is probably best handled with a digital controller where changes in gain settings can be quickly altered through swapping variable values in a program rather than switching resistor values on a circuit board.

- *Does the potential for sales volume justify a custom controller for the application*? In many high volume applications—particularly those which do not require a display or those in which a modest amount of error is tolerable—a custom controller may be an excellent choice. The controller can be designed to provide the most cost-effective solution while providing only those features which are essential. In most cases, a prototype controller can be tuned to the TE system and those settings can be translated into hard-wired gain settings for the production version of the control board. From that point onward, trimming of controller settings will be confined to those areas of the controller where it is absolutely necessary for delivering a specified level of performance (e.g., calibration of the temperature sensor). This approach can result in substantial savings. Even when there is a need for displays and high-accuracy circuitry, custom solutions can sometimes yield significant cost reductions in high-volume applications.

It should be noted in this discussion, that low-cost, stripped-down designs do not tend to be flexible; they are spec'ed for a particular situation and redirecting them to another purpose can involve significant costs and effort. For example, if you want to use a 'bare bones' controller from one application with a different thermal load that requires another set point or range, many (perhaps most) of the fixed-value resistors on the controller's circuit board will have to be changed. This will involve some mathematical analysis. When the controller is thermistor-based—and many of the lower-cost circuits will be—if you decide to specify a different thermistor, values for most resistive components will have to be recalculated. Certainly, it is possible to create flexible architectures for custom controllers, but the greater the adaptability demanded, the more likely it is that costs will soar. At some point of complexity, the hardware and labor costs will commend going with an off-the-shelf solution that has ‘all the bells and whistles’.

- *What sort of temperature sensors should be used*? This is a very important question and must be considered within the overall objectives for TE system performance. For example, in a single set-point application with no display, or in a situation where a ‘blind’ set-point range can be used, a simple thermistor offers a number of cost and utility advantages. With a thermistor, the designer can maximize the volts per degree ratio to a level which can offer much more noise immunity than is possible with many other types of sensors—this alone can help to eliminate a significant number of components from the controller. If a digital readout or scaled set-point graduation is required, integrated circuit sensors (e.g., LM324, LM335, or AD590) may be the best choice because they offer good linearity and relatively high accuracy after calibration (although their relatively high mass rules them out in applications which require greater sensitivity). Many designers prefer to use thermocouples in their applications, but the need for cold-junction compensation and low-noise circuitry will add significant expense—thermocouple-based solutions never come cheaply. Of course, off-the-shelf digital controllers will almost always come equipped with provisions for thermocouples.

- *Is ramp rate an important consideration*? Many designers seek to maximize the ramp rate in reaching the set-point and they need to understand that there are a number of important related issues. Foremost among them, is the reality that excessive ramp rates can actually destroy TE modules. Because of different rates of thermal expansion and contraction in the materials which make up TE modules, excessively fast temperature changes can cause solder joints to open—and a single open solder joint renders the module useless. This issue is especially important when the technology is used in heating applications, because TE modules are more efficient as heaters than coolers. The absolute maximum ‘safe’ ramp rate, is approximately 1° C per second and it is recommended that the end temperature be held for several minutes in such cases. Thus, whenever TE modules are used in a system in which excessive ramp rates are possible, a suitable temperature controller must be employed to limit the rate of change to a safe level. In some systems, this can be accomplished through appropriate adjustment of a derivative amplifier, in others a ‘ramp and soak’ controller will be necessary (these latter controllers ramp the set-point up or down to fit a programmed temperature profile).

- *What is the nature of available power for the controller*? If AC utility power is readily available, the user has many more options and has a better chance to keep costs down. Most off-the-shelf control solutions run on 120 or 240 VAC (although you will still need a DC supply for the TE module). If the only source of electrical power is DC (e.g., the 13.2 VDC in most automotive systems), there will be far fewer options—in many cases it will be necessary to go with a custom solution.

- *If steady-state operation is desired, can pulse-width modulation be employed*? In the vast majority of applications, pulse-width modulation at 60 Hz or greater can be used with relative impunity as long as the designer is careful to keep signal wiring separate from power wiring, uses suitable shielding with each, and follows good circuit board design practices. When done properly, control with pulse-width modulation can be employed in a far more cost-effective manner than linear control. On the other hand, if the user is dealing with very sensitive circuitry—particularly if it is to run from the same power supply as the TE system—it may be necessary to use linear drive; this is often the case in the laser diode world, for instance. When making borderline calls, especially in low-volume applications where the user cannot afford the time to experiment with control options, it can make sense to go with linear control from the beginning—there's no sense flirting with EMI problems if you don't have time to address them. In the vast majority of applications, however, PWM will work just fine and it can be a costly mistake—especially in high volume applications—to simply assume that linear control will be required. Again, PWM of thermoelectrics *has* been used successfully in automotive applications where the final product was subjected to demanding EMI/EMC testing.

Tellurex Temperature Control Products

When it comes to temperature control, while Tellurex Corporation offers a number of alternatives, it truly specializes in providing solutions for high-volume applications where a low-cost, custom controller is suitable. In most cases, we can adapt a previous design to new systems and then provide circuit boards which are custom populated. In those cases which require unique approaches, we can rely upon our extensive experience to arrive at optimal solutions. To discuss your needs, contact a sales representative at Tellurex (231-947-0110 or sales@tellurex.com).

In Conclusion

While TE technology is very controllable, there are a great many things to consider in arriving at the best solution. Every desire comes at a cost, so each must be contemplated thoroughly before proceeding. The more that cost is a constraining force in choosing a design path, the more diligent the designer must be in reigning in the ‘wish list’ to focus on what is absolutely necessary. The staff at Tellurex is ready to assist you in exploring options.