Thermoelectric Temperature Control Using the ispPAC20

September 2001 Application Note AN6029

Overview

Temperature controls are found in many places, ranging from steel mills to the thermostat in your living room. This application note focuses on one particular type of temperature control system, based on the use of Peltier-effect thermoelectric coolers. Temperature control systems based on this technology are useful in situations where the temperature of a small object must be maintained with a high degree of precision, such as laser diodes in fiber optic telecommunication systems.

This application note is organized into topics according to the following outline:

1. Thermoelectric Coolers
2. Temperature Control
3. Driving a TEC from an ispPAC20
    - 3.1 The H-Bridge Driver
    - 3.2 Current-Output H-Bridge
    - 3.3 Feedback and Linearization
    - 3.4 Gain Stabilization
    - 3.5 Current Limiting
4. Using Thermistors as Temperature Sensors
5. Feedback Control Systems
    - 5.1 Proportional Controller
    - 5.2 Dominant Pole Controller
    - 5.3 Integral Controller
    - 5.4 Proportional-Integral (PI) Controller
6. Example System

1. Thermoelectric Coolers

When designing a small temperature control system, the best choice for a refrigerator/heater element is often a thermoelectric cooler (TEC). The operation of these devices is based on the Peltier effect, which was discovered in 1834 by Jean Peltier. When an electrical current is passed thorough the junction of two dissimilar metals (Figure 1), heat is also transferred across the junction.

Figure 1. Peltier Effect across a Bimetallic Junction

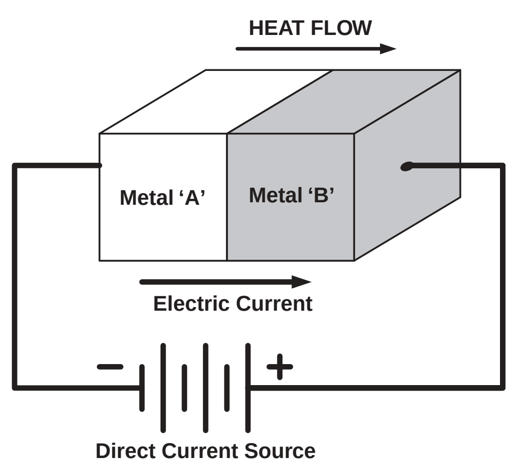

It is difficult to make a practical cooler out of most common metals, however, because the thermal conductivity of metals is very large in relation to the magnitude of their respective Peltier coefficients. This results in a small amount of heat being transported, which rapidly diffuses back across the junction. Some semiconductors, however, most notably Bismuth Telluride, have a very pronounced Peltier effect, and low thermal conductivity. The combination of being able to efficiently pump large amounts of heat and not have it diffuse back across the cooler allows for the realization of practical refrigeration devices. Figure 2 shows a schematic view of a typical thermoelectric cooling module, constructed from alternating pellets of 'N' and 'P' type semiconductor material electrically connected in series. When direct current is applied to this assembly, heat is removed from the top surface of the cooler module and is transported to the bottom surface, cooling the top surface and any attached assembly. If the polarity of the applied current is reversed, the cooler pumps heat in the opposite direction, heating the top assembly. This feature of a Peltier effect cooler is extremely useful, in that it allows one to build temperature control systems that are capable of both heating and cooling their thermal loads to maintain proper temperature.

Figure 2. Practical Thermoelectric Cooler Assembly

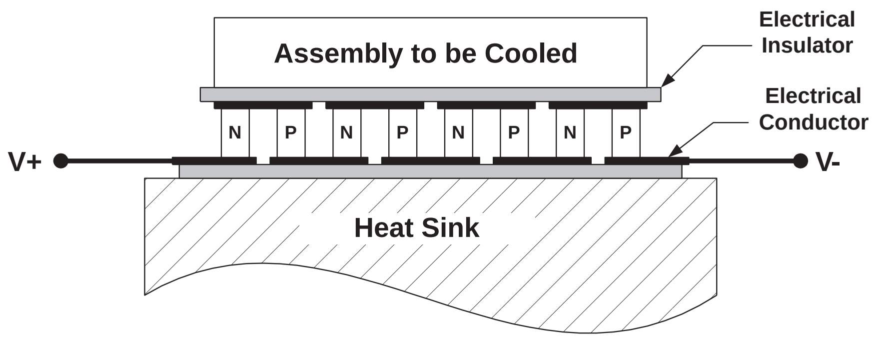

2. Temperature Control Loops

If one needs to maintain a stable temperature, one of the best ways is to use a feedback control system. In a feedback control system, one measures the actual temperature of the object of interest, compares that measurement to a desired *setpoint* temperature, and then appropriately heats or cools the object to try to bring it towards the setpoint. Figure 3 shows an example of a feedback temperature control system using a TEC.

Figure 3. Generalized Thermoelectric Temperature Control System

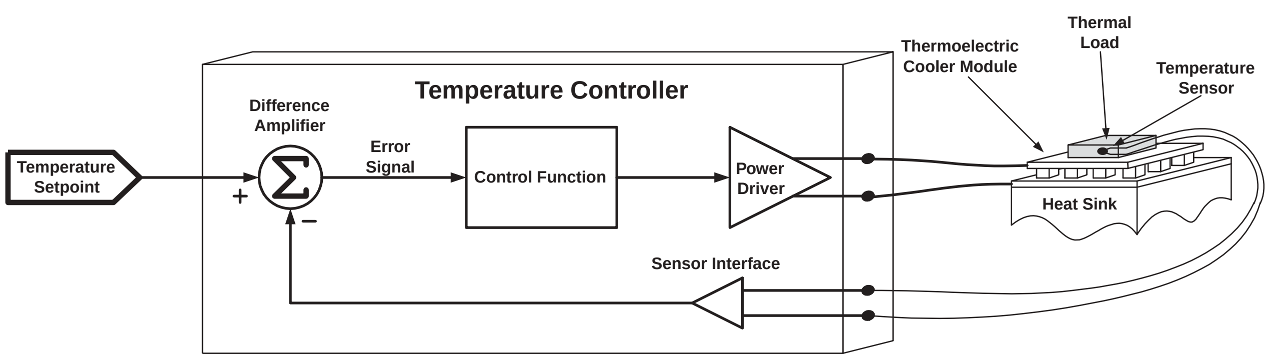

While the above system may seem simple, there are a great many implementation details that must be considered when building a practical temperature control system. A few of the more important issues are:

1. Actuation - Providing power to drive the thermoelectric cooler.
2. Sensing - Accurately measuring the thermal load's temperature.
3. Intelligence - Determining how much to drive the thermoelectric cooler to get it to settle to the setpoint.

This application note discusses each of these topics and provides example circuits for performing the required function.

3. Driving a TEC from an ispPAC®20

To pump heat, a thermoelectric cooler can require a significant amount of power, with typical drive voltages ranging from 2V to 12V and typical drive currents ranging from a fraction of an ampere to tens of amperes. Because of the amounts of power needed to drive a TEC, power-interface circuitry is necessary. Additionally, because a TEC can be used to both heat and cool its thermal load by controlling the direction of the applied voltage, it is also desirable in a precision temperature control to take advantage of this heat/cool capability.

While it is straightforward to implement a circuit that can switch a TEC from the OFF state to either FULL COOL or FULL HEAT, this approach has several disadvantages. The first, and most obvious of these is that this type of control scheme may make it difficult to precisely control temperature, especially if it takes some time for heat to diffuse through the thermal load to the sensor. The preferred way to power a TEC module for high accuracy temperature control is from a continuous voltage or current source.

3.1 The H-Bridge Driver

Many contemporary electronic systems are designed to operate from a single positive voltage supply. While single-supply operation presents no serious difficulties if one wishes to operate a TEC in a heat-only or cool-only mode, it makes heat/cool operation a little more complex. Fortunately, a circuit topology known as an H-bridge can provide bipolar drive to a TEC or other load while still operating from a single supply voltage. Figure 4 shows a simplified H-bridge driver, controlled from one of the differential outputs of an ispPAC20.

Figure 4. Basic H-Bridge Driver

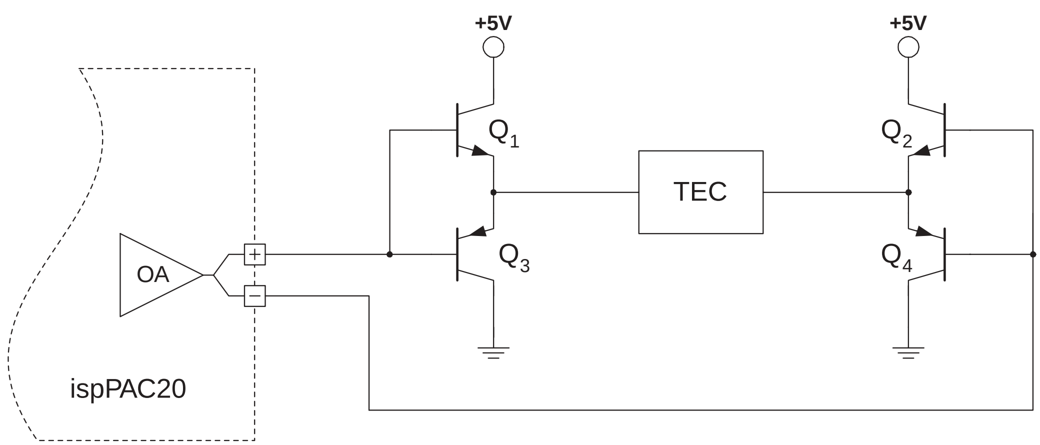

In this circuit, a positive output signal (V+ > V-) from the ispPAC20 will cause Q_1 and Q_4 to turn on in a proportional manner, causing current to flow from the +5V rail to the emitter of Q_1, through the TEC module, and into the emitter of Q_4 before returning to ground. In the case of a negative output signal from the ispPAC20, transistors Q_2 and Q_3 will turn on, and will result in a current flowing in the opposite direction through the TEC. The amount of voltage applied across the TEC, and consequently the amount of current flowing through it, will be proportional to output voltage of the ispPAC20. A key feature of this circuit is that at no time do all the transistors turn on – this would effectively short the power supply to ground and cause serious problems.

As drawn, this circuit is capable of driving small TECs with some limitations on operating voltage and current. First, because the transistors are being used in emitter-follower configurations, there will be ≈0.7V drop from the base to the emitter of each active device. If the ispPAC20 provides a differential output of 4V, only 2.8V (4V - 2 (0.7V) can be expected to be seen across the TEC. This voltage drop will limit the maximum output voltage swing to considerably less than the 5V available at the supply.

The second limitation of this circuit is that of limited output current, stemming from the finite DC current gain (H_{FE}) of the transistors used in the bridge. Because the ispPAC20 can only sink or source 10mA of output current, the maximum current available to drive the TEC is limited to H_{FE} x 10mA. For a robust design, one will want to specify a transistor based on its minimum H_{FE}. Typical minimum H_{FE}'s for small power transistors range from 30 to 100 and will therefore limit the maximum available current to drive the TEC to 300-1000mA.

One final, and subtle consideration in this design is that of stability. As drawn, the H-bridge circuit of Figure 4 may oscillate at RF frequencies, depending on the characteristics of the transistors used. To prevent spurious high-frequency oscillation it may be necessary to add 'lossy' components such as resistors or ferrite beads in series with the base or emitter leads of the transistors.

3.2 Current-Output H-Bridge Driver

When higher currents or voltages are required to drive the TEC, a different output stage may be needed. In the circuit of Figure 5, two stages are used to provide high-current drive capabilities.

Figure 5. H-Bridge Power Driver for ispPAC20

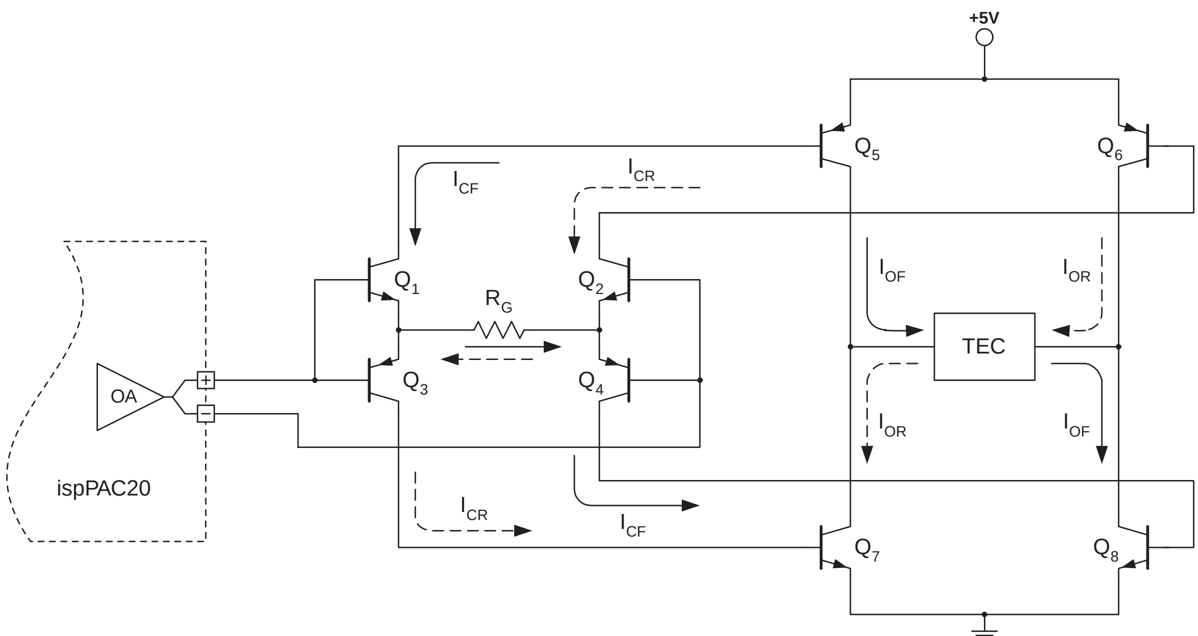

This circuit works in the following way: When the differential voltage at the output of the ispPAC20 swings positive (Vout+ > Vout-), transistors Q_1 and Q_4 begin to conduct. This causes a voltage to be forced across R_G, resulting in current I_{CF} to flow through Q_1 and Q_4. I_{CF} flows into the bases of power transistors Q_5 and Q_8, where it is amplified to I_{OF}, powering the TEC in a positive direction. In the case of a negative voltage at the ispPAC20's output, Q_2 and Q_3 are turned on, resulting in a current of the opposite polarity (I_{CR}) to flow through R_G. I_{CR} is amplified by power transistors Q_6 and Q_7, biasing the TEC in a negative direction.

Because none of the transistors are switched on hard into saturation, and remain in linear mode, the TEC drive current will be proportional to the voltage appearing at the ispPAC output. Figure 6 shows the DC transfer function. Note that the transfer function for this circuit includes a 'dead-zone' in which small input voltages result in zero output drive current. This dead-zone is important because it ensures that one set (e.g. Q_5, Q_8) of power transistors will be completely off before the other set (Q_6, Q_7) begins to turn on. Allowing both sets of power transistors to turn on simultaneously is a bad thing, as it allows current to flow directly from the positive supply to ground, bypassing the TEC and causing excessive heating in the output transistors.

In addition to being able to handle more current, this circuit is also different in several other ways from the circuit of Figure 4. The first difference is that this circuit provides rail-to-rail differential drive, and can apply a voltage of nearly ±5V across the TEC when operated form a +5V supply.

Another difference is that where the circuit of Figure 4 applies a voltage to the TEC, this circuit applies a current. Current-mode output provides two advantages over a voltage mode output. The first is that it offers a degree of protection to the TEC, as circuit component values can readily be chosen to ensure that a given maximum output current is never exceeded. Another more subtle advantage is that because the heat pumped by a TEC is proportional to current, current-mode control allows for more uniform operation of the TEC than does voltage-mode control.

Finally, because signals are transmitted in the form of current, it is possible to operate the power output stage of this circuit from a different set of rails than the ispPAC20. For example, if a TEC were to require a ±12V drive signal, the output bridge (Q5-Q8) could be operated from +12V and ground. Operation from a lower voltage supply (e.g. 3.3V) is also possible, but may require that the common-mode output voltage of the ispPAC20 be adjusted (via the CMVIN pin) to center the devices output voltage between the supplies.

Figure 6. DC Transfer Function of H-Bridge Driver

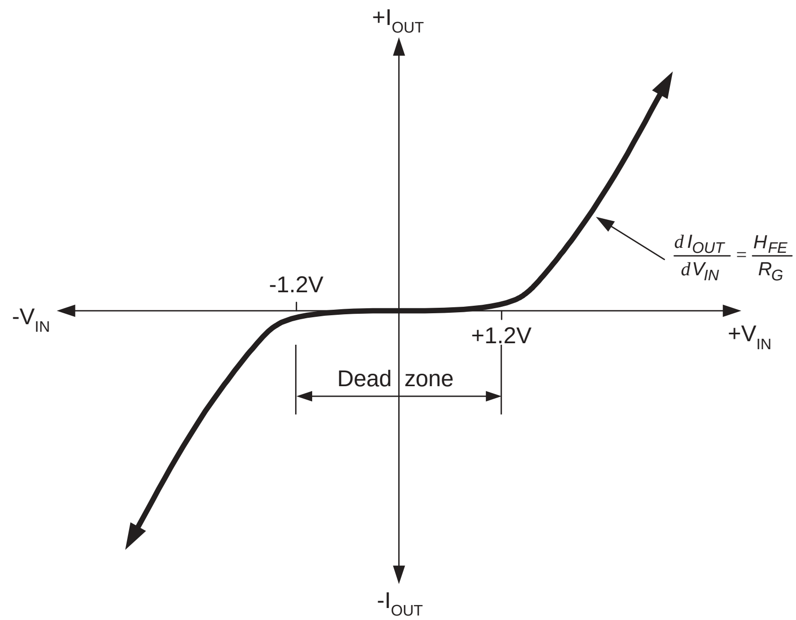

3.3 Feedback and Linearization

While the presence of the dead-zone is a good thing for proper circuit operation, it is not as desirable a feature from the standpoint of building a linear control system. Non-linearities such as the dead-zone can reduce the system's stability and accuracy. Non-linear functions in control loops also can make the systems more difficult to analyze and understand, and therefore less predictable. For this reason it would be nice to be able to reduce or eliminate the dead-zone, while still ensuring that the output transistors do not all turn on at once. One way to do this is to use a feedback loop, as shown in Figure 7.

Figure 7. Using Feedback to Linearize Response Function

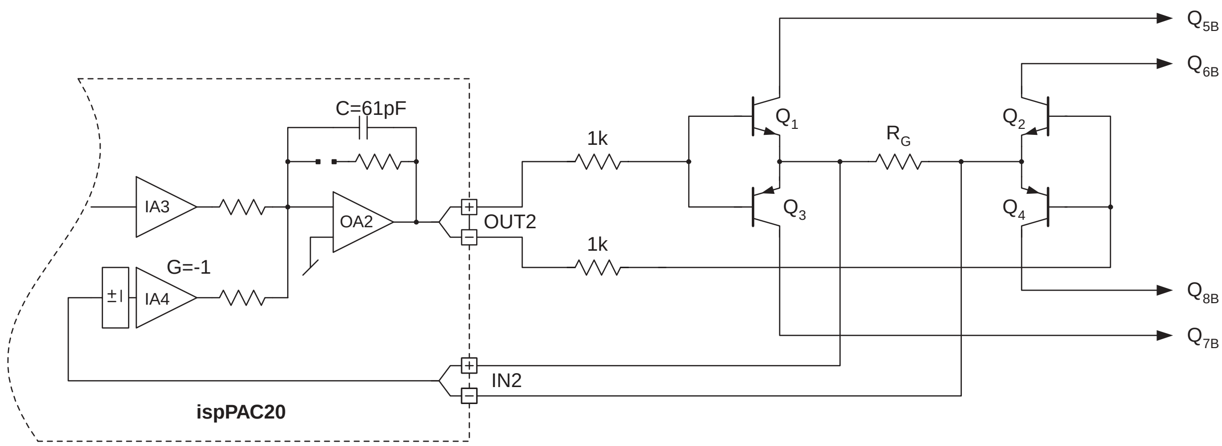

This circuit senses the voltage across R_G, which is proportional to the current flowing through this resistor. A negative gain for input amplifier IA4 creates a negative feedback loop. This feedback loop will try to swing the output of output amplifier OA2 to maintain the condition

$$\left(\frac{V_{RG}}{V_{IN}}\right) = -\left(\frac{G_{IA3}}{G_{IA4}}\right) \tag{1}$$

If a small positive voltage is input to IA3, the feedback loop will cause the OA's output to swing sufficiently positive to get past the driver's dead-zone. Conversely, a small negative input will result in the OA's output swinging sufficiently negative to get out of the driver's dead-zone. This results in the linear transfer function shown in Figure 8.

Figure 8. Linearized Driver Response Function

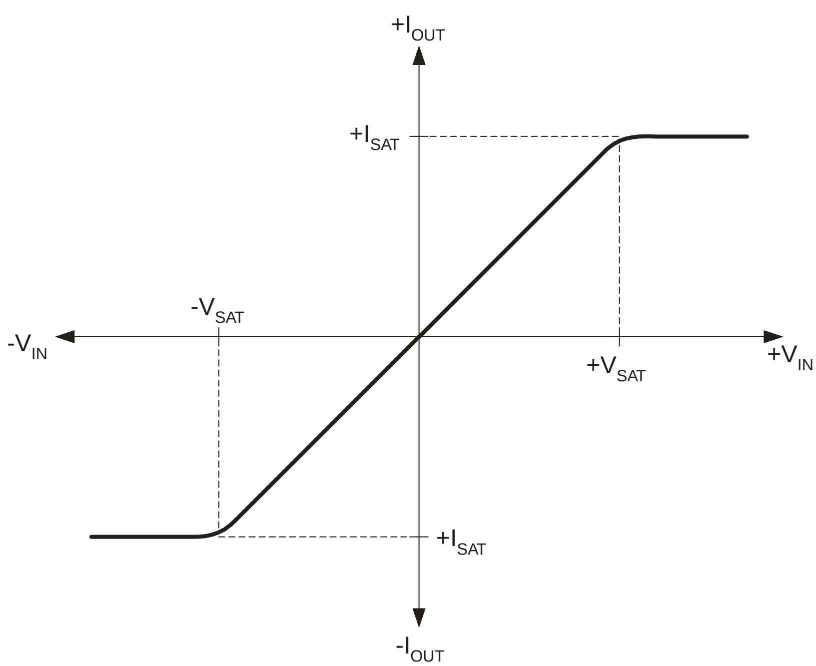

By using a fast negative feedback loop, it is possible to linearize a highly non-linear circuit. The linear relationship between input voltage and output current for this circuit can now be approximated by a constant driver gain A_D

$$A_D = \frac{I_{TEC}}{V_{IN}} = -\left(\frac{G_{IA3}}{G_{IA4}}\right) \times \left(\frac{H_{FF}}{R_G}\right) \qquad (2)$$

where H_{FE} is the gain of the power transistors used in the H-Bridge driver (Q_5, Q_6, Q_7, Q_8).

3.4 Gain Stabilization

H_{FE} represents the DC current gain of a transistor, or the ratio between base current and collector current. For most production transistors, this parameter can vary over a factor of 4 or 5 or even more from unit to unit, and can also vary considerably for any given device over bias conditions and temperature. Because the gain of the output driver stage is a direct function of H_{FE}, this circuit's gain will also see a significant spread over manufacturing variation and operating conditions. While this may not be a problem for many controllers, such variations can be an issue in the case of a highly engineered, tightly tuned control loop. For this reason it may be desirable to reduce effects of H_{FE} variation on system performance. One way to do this is by stabilizing the gain of the transistors in the output driver. Figure 9 shows an H-bridge with gain-stabilized output transistors. In this circuit, scaled current mirrors are used to control gain. The easiest way to understand how this works is through an example. To get a gain of ~25, one might set R_1=1Ω and R_2=R_3=25Ω. A current of 20mA from the collector of Q_3 will flow through Q_9 and establish a voltage of ~500mV (20mA x 25Ω) at its emitter, and approximately 1.2V at its base. This in turn will result in ~500mV at the emitter of Q_7, and a corresponding current of 500mA (500mV/1Ω) through R_1. The relationship between input current (I_{CQ9}) and output current (I_{CQ8}) over a wide range of current levels can be approximated by:

$$\frac{I_{CQ8}}{I_{CQ9}} \approx \frac{R_2}{R_1} \qquad (3)$$

The major assumptions behind this approximation are that the voltage drops across the emitter resistors (R_1, R_2, R_3) are comparable to the transistor base-emitter voltage drop (~0.6V) and that the H_{FE} of the power transistors is significantly greater than the desired current gain for the H-bridge. It is not possible with this circuit to get a current gain that exceeds the H_{FE} of the output driver transistors, as it sacrifices gain to obtain predictability. Because it is necessary to drop some voltage across the emitter resistors, this circuit also sacrifices some output voltage compliance range (0.5-1V) as compared the circuit of Figure 5.

Figure 9. Stabilizing the Driver Gain against Variations in HFE

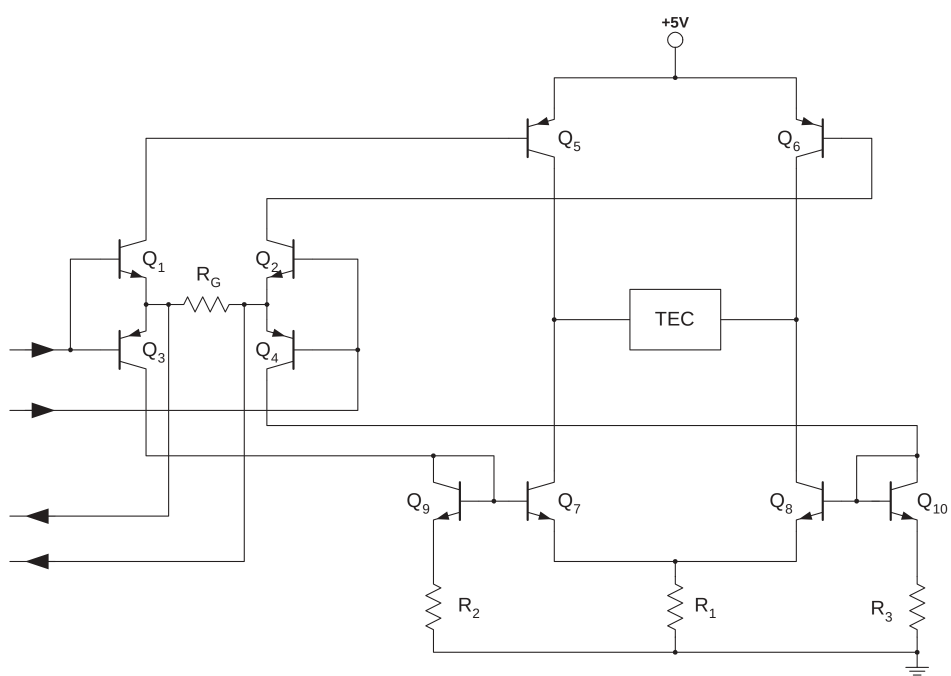

3.5 Current Limiting

Another issue relating to TEC driver circuits is that of current limiting. Because the driver circuits presented here provide their output in the form of a controlled current, as opposed to merely impressing a voltage across the TEC, they are inherently less likely to damage the TEC by an overcurrent condition than a voltage-mode output driver would be. In those cases where one cannot or does not wish to rely on the controller to not overdrive the TEC, it is possible to build hardware current limiting directly into the H-bridge driver. The advantage of implementing current limiting in this manner is that regardless of what the controller does, it will not be able to overdrive the TEC. Figure 10 shows an output driver incorporating a simple ‘foldback’ current limiting scheme.

Figure 10. Foldback Current Limiting Circuit

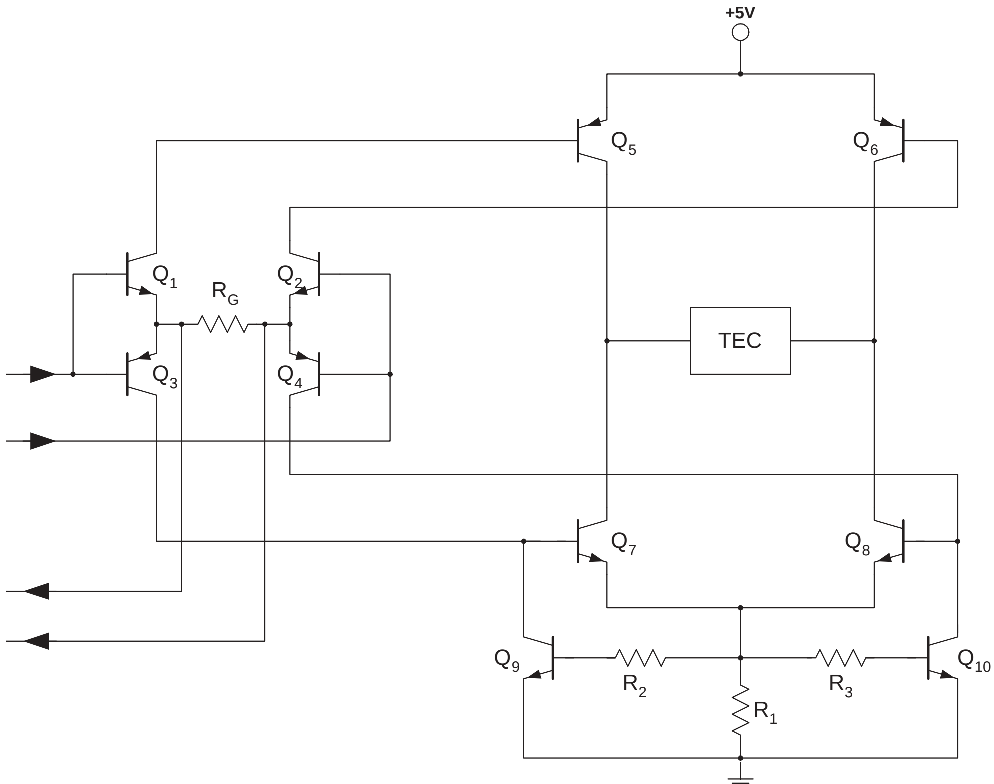

Again, the easiest way to understand how this circuit operates is through an example. If Q_7 is conducting, its output current also passes through R_1, where it results in a voltage. If this voltage is less than 0.6V, Q_9 remains off. If the voltage across R_1 increases to ~0.6V, it begins to turn on Q_9 (through R_2), which steals current from the base of Q_7, and tries to turn it off. Because this process is continuous, the overall result is that the current through Q_7 will be limited to a maximum level I_{MAX} than can be approximated

$$I_{MAX} \approx \frac{0.6V}{R_1} \qquad (4)$$

Because the circuit is symmetric, and only one half is active at any time, the same behavior holds true for the leg comprising Q_8, Q_{10}, R_1 and R_3.

The purpose of R_2 (and R_3) in this circuit may not be immediately obvious, but it is to improve the circuit’s high-frequency stability, and will typically be in the range of 10-100Ω. Alternatively, a ferrite bead may also be useful for this function, and will not affect the DC performance as a resistor does (increases limit current).

4. Using Thermistors as Temperature Sensors

There are many types of temperature sensors that can be used in a temperature control system. Two of the major considerations in selecting an appropriate sensor are the temperature range over which it is expected to function and its accuracy. For applications involving moderate temperatures and where accuracy on the order of 1°C is required, thermistors are often a good choice for a temperature sensor. A thermistor is a slug of material whose electrical resistance varies as a function of temperature. While the electrical resistivity of conductive materials (metals, semiconductors) is temperature dependent, the materials used in constructing thermistors typically exhibit large and well characterized temperature coefficients of resistivity. As an example, the graph of Figure 11 shows the resistance of a 10kΩ type 'J' thermistor over a temperature range of -20°C to +70°C. Although the thermistor's resistance changes considerably over this range, the response is not linear with respect to temperature.

Figure 11. Resistance vs. Temperature for 10kΩ Type 'J' Thermistor

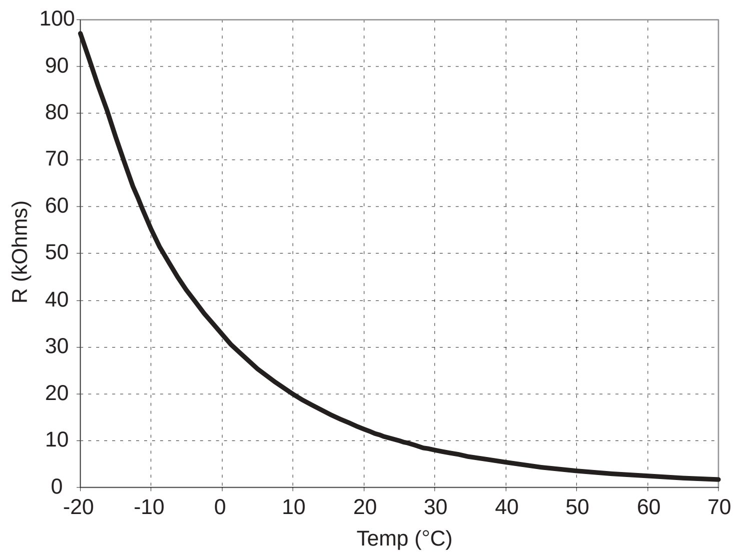

To be useful as a temperature sensor, it is necessary to measure the thermistor's resistance. While there are many ways to do this, one of the simplest is to incorporate the thermistor in a voltage divider or balanced bridge circuit, in which the thermistor's resistance is compared to a known fixed resistance. Figure 12a shows how a voltage divider can be used to generate a single-ended temperature dependent voltage. At the thermistor's nominal resistance temperature (often 25°C), this circuit will generate a 2.5V output voltage. As temperature increases, the voltage will drop.

Because the ispPAC20 has differential inputs, the differential balanced bridge circuit of Figure 12b will often be a better way to bias the thermistor. In this case, the bridge output voltage will be zero at the thermistor's nominal resistance temperature. For situations in which the normal operating range of the thermistor is significantly different than its nominal temperature, one can re-scale the other bridge resistors to match, providing a zero output voltage at this other temperature. For example, if one wished to use this thermistor to measure temperatures over a range of –20°C to 0°C, one might use 54.9KΩ resistors in the bridge to match the 55.3kΩ resistance of the thermistor at –10°C (the center of the desired measurement range). In this case, the zero-voltage output would occur when the resistance of the thermistor matched the resistance of the other resistors in the bridge (≈-10°C).

Figure 12. Single-ended and Differential Thermistor Interface Circuits

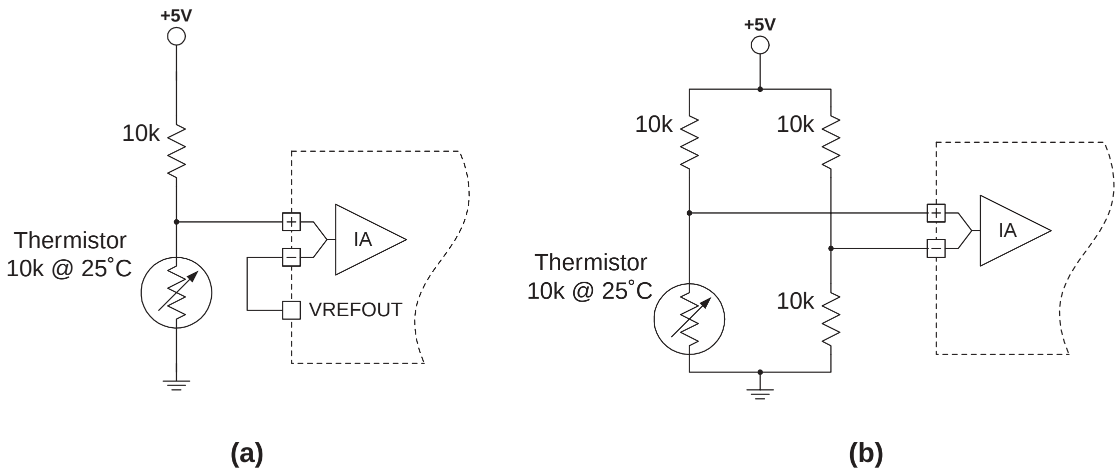

An additional benefit of using the interface circuits of Figure 12 is that their voltage output as a function of temperature is significantly more linear than the thermistor's resistance vs. temperature characteristics. Figure 13 shows the voltage vs. temperature characteristics of the bridge circuit (Figure 12b) when using a 10kΩ type 'J' thermistor. Note that over the temperature range surrounding the thermistor's nominal value temperature, the response is fairly linear.

Figure 13. Thermistor Bridge Output Voltage vs. Temperature

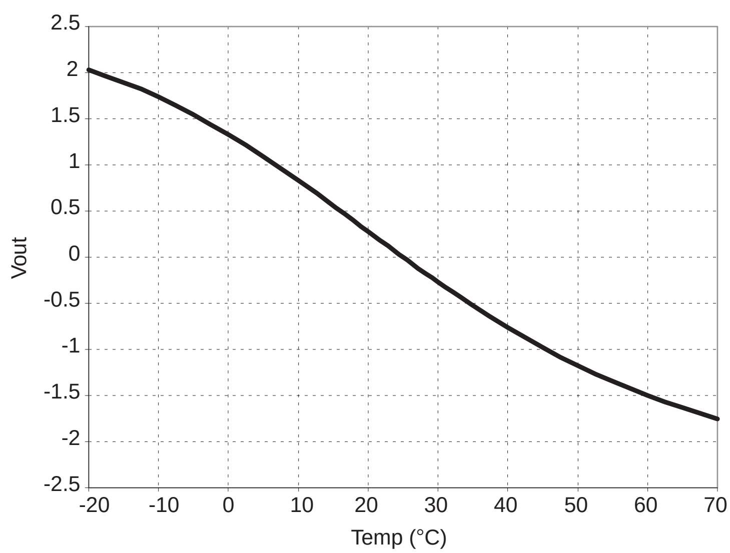

5. Feedback Control Systems

So far, we have examined how to provide power drive functions for a TEC module and how to measure temperature using a thermistor. The remaining piece we need to implement a temperature control system is the controller itself. Before jumping into a description of specific types of controllers that can be used, it is worth discussing a few of the characteristics and design issues involved in successfully implementing a control system.

Figure 14 shows a block diagram of a generic linear feedback control system. A linear feedback control system works by comparing a measurement of the controlled process variable to a desired setpoint, by subtracting the two signals. This results in an error signal, which indicates the degree and direction (higher or lower) to which the setpoint and the process variable mismatch. This error signal is then transformed by a control function into a control variable, which determines how much to drive the controlled system.

Figure 14. Classical Feedback Control System

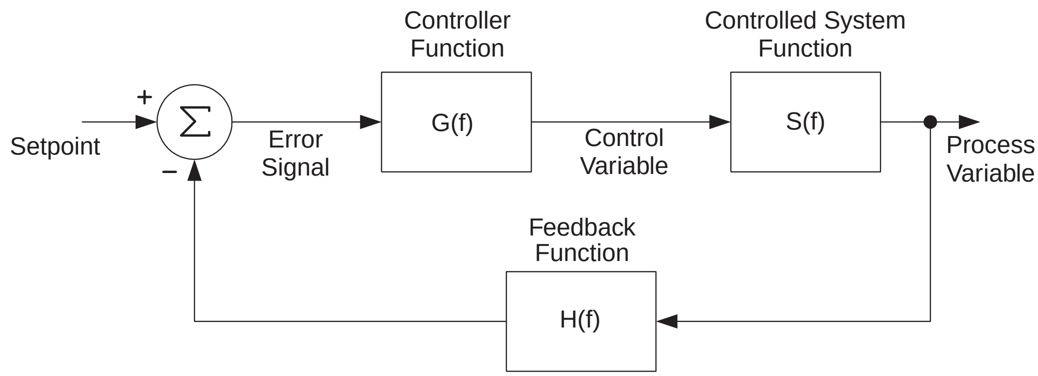

One of the factors that makes control system design challenging is that the controlled system and feedback functions may have time-dependent behavior. In the case of a temperature control, turning on a heater does not make the object being heated instantaneously rise to a desired temperature; there can be a significant time lag. Similarly, the thermocouple used to measure the system's temperature can also have a noticeable response time. These time lags can make the system difficult to control well.

Two of the major concerns in a control system are stability and settling time. A stable control system will eventually bring the process variable to a steady-state condition within a given error bound. An unstable system will either force the process variable to some limiting high or low condition, or it will oscillate between two or more states indefinitely. In the case that a system is stable, and does converge to a low-error steady-state condition, the amount of time needed to do so is the settling time. In most control systems, fast settling time is a desirable feature.

Although the subjects of how to achieve stability and fast settling times are very complex, there are some relatively simple, rule-of-thumb techniques for quickly and approximately characterizing a linear control system. One such technique is that of determining loop gain and phase margins.

If one assumes that each of the loop functions is linear (or close enough for practical purposes), one can characterize it by its gain and phase responses. If one 'unrolls' the control loop into a linear chain, as shown in Figure 15, one can then easily determine the gain and phase response for the whole chain. (Hint: if the gain is characterized in dB, and phase in degrees, one can obtain the total gain by just adding all the separate dB curves and obtain the total phase shift by just adding all the separate phase shifts.

Figure 15. Unrolling the Control Loop

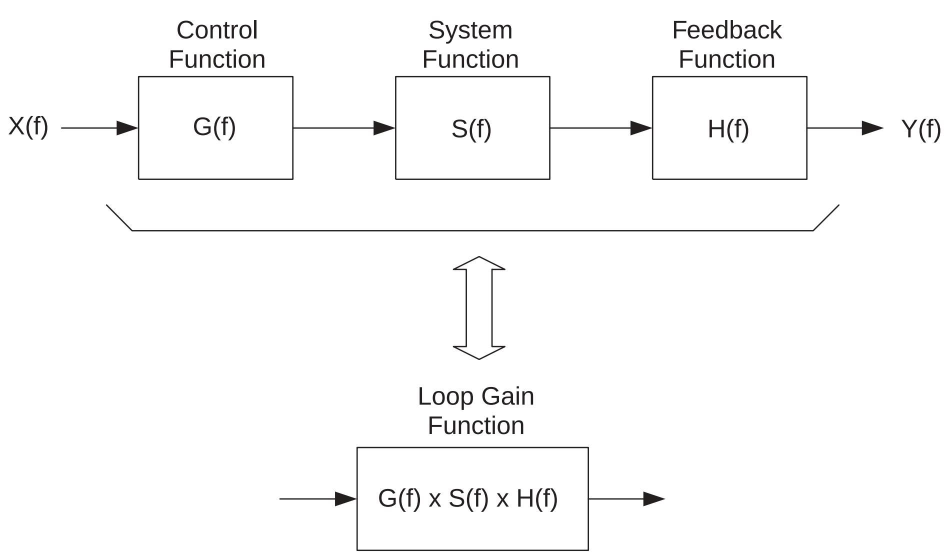

By looking at the chain response, one can quickly get an idea of whether the system will be stable or unstable. Figure 16 shows example system gain and phase response curves. These curves may be more familiar to electronic designers when referred to as Bode plots.

Figure 16. Response of Unrolled Loop, showing Gain and Phase Margins

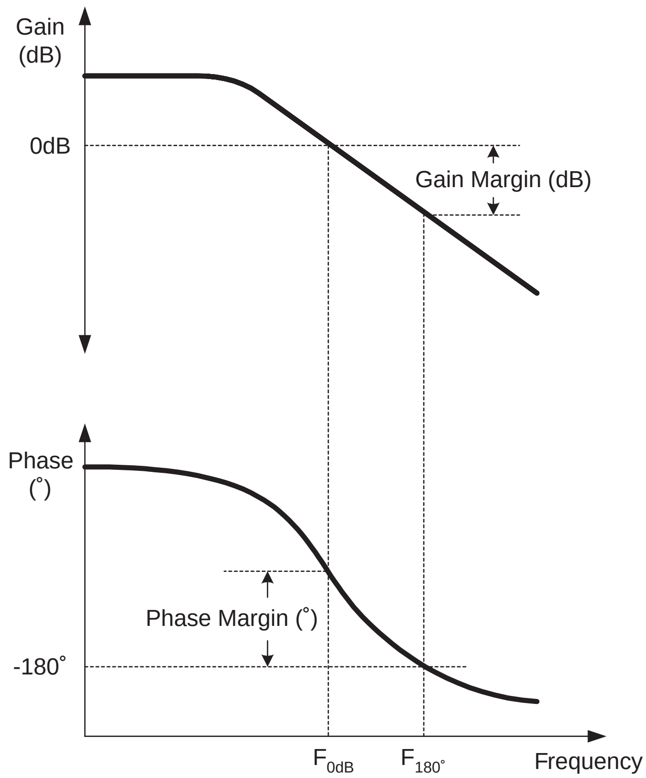

Stability in a linear control system depends on maintaining negative feedback, which is accomplished by subtracting some measured version of the process variable from the setpoint. If one instead *adds* the process variable to the setpoint, this is known as positive feedback, and will result in an unstable condition in which the controller will drive the process variable to some minimum or maximum limiting condition.

In a system with a sufficient amount of time lag around the control loop, the negative feedback can become positive feedback at some frequency. This point occurs when the phase lag around the loop is 180°. If the total loop gain at this frequency is even slightly greater than one, the loop will begin to oscillate at this frequency. The two main indicators of stability, based on the gain and phase response of the control loop are called the *gain margin* and *phase margin*.

A control loop's gain margin is the amount by which the loop's gain is less than one at the frequency where the phase lag is 180° (F_{180°). An ideal linear system with positive gain margin (gain < 1 at F_{180°) will be stable. One which has zero or negative gain margin (gain >=1 at F_{180°) will be unstable. Gain margin is usually expressed in dB.

Similarly, one can characterize a system's stability by how close one is to 180° of phase shift when the gain is one. This is called the phase margin and is expressed in degrees. A system with a positive phase margin will be stable, while one with a zero or negative phase margin will not.

Even when one has a stable system, the gain and phase margin can affect the system's qualitative behavior (Figure 17). A system with very little phase margin may exhibit an underdamped behavior in response to changes in setpoint, where the value of the process variable will overshoot the setpoint and 'ring' or 'seek' as it settles to a final value. In the opposite case, a system with a great deal of phase margin may exhibit an overdamped response, in which the process variable does not overshoot the setpoint, but takes an unacceptable amount of time to settle. For most systems, one will want to try to tune the controller to achieve something close to a critically damped response, in which the system settles quickly, but without overshooting the setpoint.

Figure 17. Under, Over and Critical Damped Responses

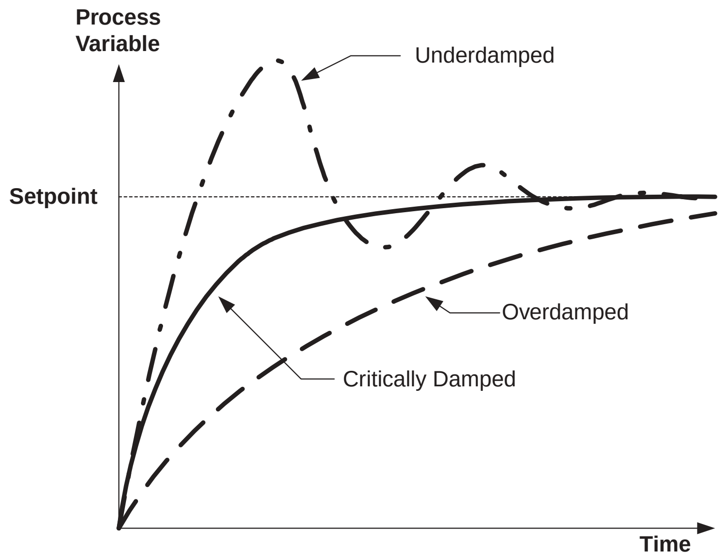

When trying to analyze control systems using techniques from linear control theory, one should keep in mind that the resulting analyses are only strictly true for ideal linear systems. Real-world systems have non-ideal characteristics which can make their actual behaviors deviate significantly from those predicted. Building an accurate model of a dynamic system is a non-trivial exercise. For these reasons, the ability to quickly and easily explore various solutions is a valuable tool when embarking on a new design. Programmable analog circuits allow the system designer to quickly evaluate numerous potential solutions.

One disadvantage, however, of having a large number of potential design solutions is that of which one to pick. Fortunately for the designer, there are several simple and relatively intuitive types of controllers that can be used as starting points in a design. The following sections will describe a few of these controllers and how they can be implemented with an ispPAC20.

5.1 Proportional Control

One of the simplest linear control strategies is proportional control. In proportional control one simply amplifies the error signal, and uses the result as the control variable. Figure 18 shows a block diagram and response characteristics of a proportional controller.

Figure 18. Proportional Controller (a) and Frequency/Phase Characteristics (b)

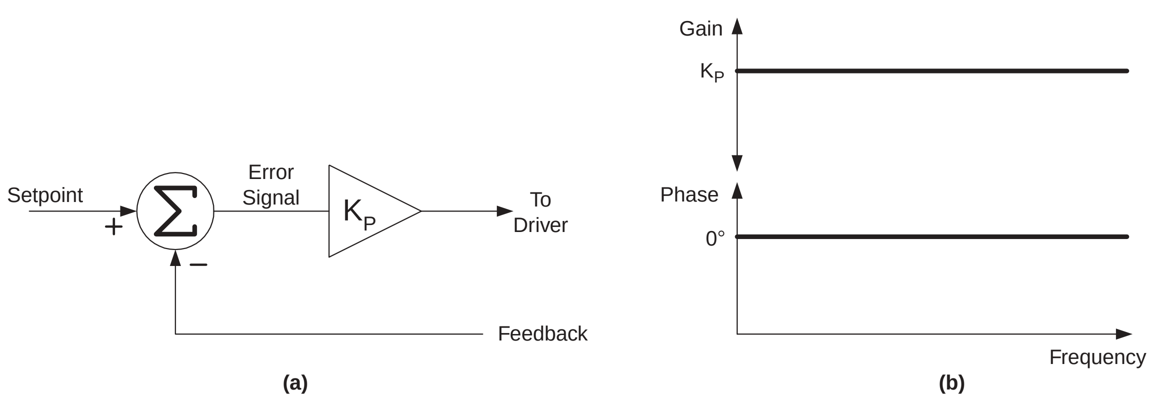

Because the magnitude of the drive signal is an amplified replica of the error signal, it follows that in order to get a drive signal, one needs a non-zero error signal. In a temperature control system, the main consequence of this characteristic is that there will be a difference between the setpoint temperature and the measured temperature of the thermal mass, and the difference will be proportional to the output of the driver required to maintain thermal equilibrium.

One obvious way to minimize this source of error is to increase the controller gain. For some kinds of systems, a proportional controller's gain can be raised to infinity without making the system unstable, at least for the case of an ideal controller and ideal controlled process. For the more commonly encountered cases, where the controller or the system is non-ideal, or the system has some second-order dynamics, there are limits to how high one can raise the gain while maintaining stability. Despite having the drawback of a finite error, a proportional controller enjoys the advantages of being easy to implement, and easy to adjust or 'tune', as there is only one parameter (K_P) that can be, or needs to be, adjusted.

There are several ways in which a proportional controller can be implemented with the ispPAC20, two of which will be outlined here. The first method (Figure 19) is applicable when one only needs a small amount of gain in addition to that provided by the TEC driver circuit. Here, the setpoint is provided by the DAC and fed into IA3, and the thermistor bridge voltage is subtracted from this signal by IA4. The error signal appears at the output of OA2.

Figure 19. Low-Gain ispPAC20 Proportional Controller

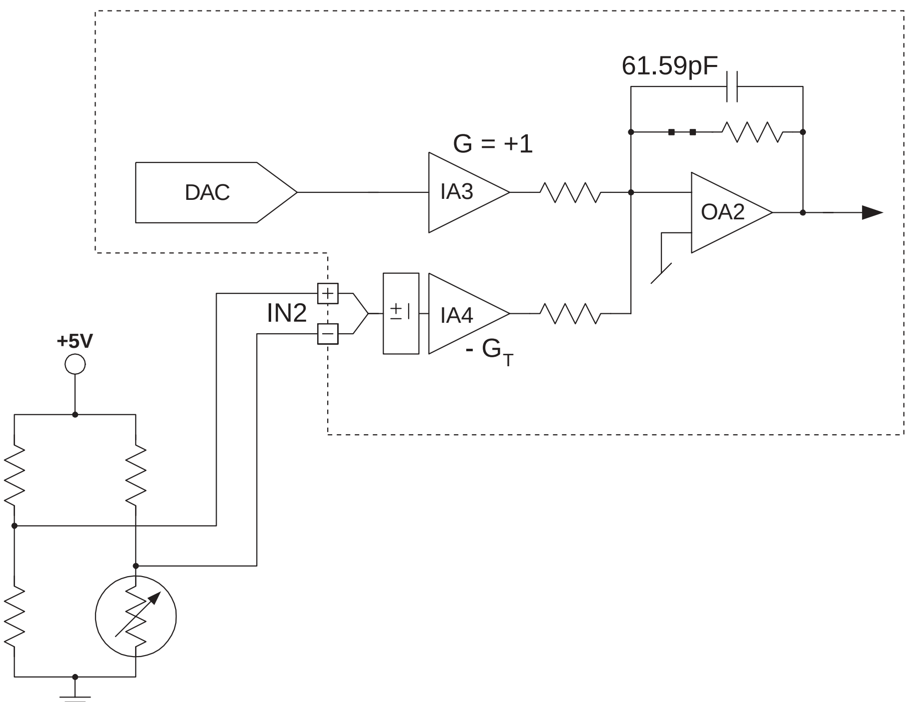

To achieve maximum resolution, one may want to amplify the thermistor signal by some value other than -1, to match it to the ±3V voltage range obtainable from the DAC. As an example, if the temperature range of interest is 15°C to 35°C, and the output of the thermistor bridge is ± 0.5V, then setting IA4 to a gain of –6 will scale that temperature range to that of the DAC voltage range. An additional effect is that providing increased gain in the sensor feedback path also increases the total loop gain, which in turn reduces the error but can also decrease the stability.

In cases where larger controller gains are required, these can be achieved by using external resistors as shown in Figure 20.

Figure 20. High-Gain ispPAC20 Proportional Controller

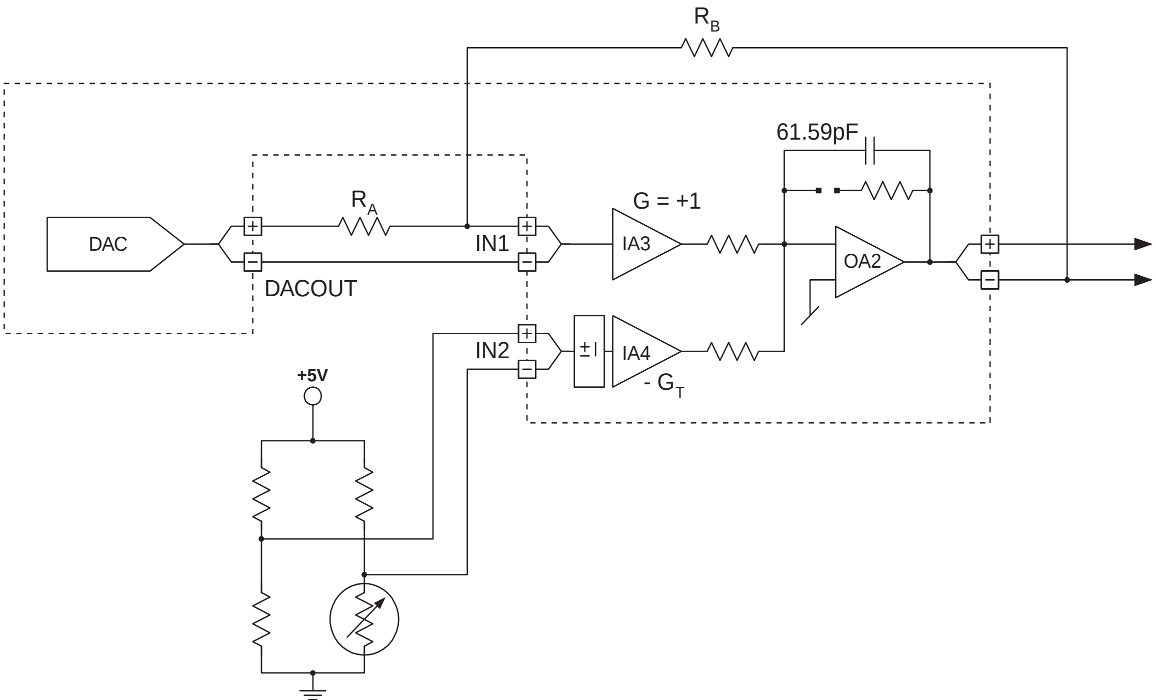

The gain for this circuit varies depending on whether it is measured from the setpoint DAC to the output or from the thermistor input to the output. From the setpoint DAC, the gain is given by

$$A_{VDAC} = \frac{2R_2}{R_1} + 1 \tag{5}$$

while the gain measured from the thermistor input to the output is given by

$$A_{VT} = G_T\left(\frac{2R_2}{R_1} + 2\right) \tag{6}$$

Varying the gain of IA4 (G_T), allows one to be able to match the input voltage range of the thermistor to the DAC output voltage range. Note that even when GT=-1, the gain from the DAC is slightly lower (-1) than that seen from the thermistor. This has the effect of modifying the scale between the DAC and thermistor.

5.2 Dominant Pole Compensation

The major drawback to proportional control is that it presents a trade-off between error and stability; increasing the gain to reduce the error also can reduce the stability. One solution to this trade-off is probably already familiar to many electronic engineers. This is the time-honored tradition of putting increasingly larger capacitors into a circuit until it stops oscillating. This technique is called *dominant pole compensation*.

The general idea behind dominant pole (DP) compensation is to reduce the gain of the controller at higher frequencies where oscillations occur. Doing so allows one to maintain a high DC gain, which reduces the steady-state error. Figure 21a shows one way of making a DP controller by putting a first order RC filter into the control signal path. The resulting controller gain and phase are shown in Figure 21b.

Figure 21. Dominant Pole Controller (a) and Frequency/Phase Characteristics (b)

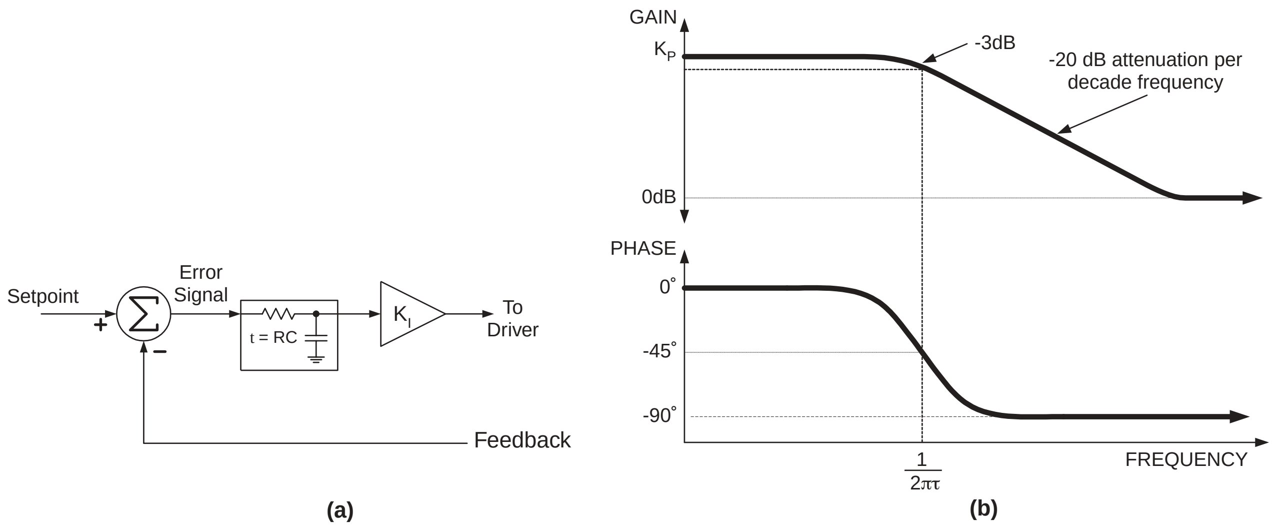

If one wants to add DP compensation to the low-gain proportional controller (Figure 19), a passive RC lowpass filter could simply be added to the controller's output. Because DP pole compensation is most useful when high DC gain is needed, a more common (and useful) scenario is when it is incorporated in a high-gain controller, as shown in Figure 22.

Figure 22. ispPAC20 Dominant Pole Controller

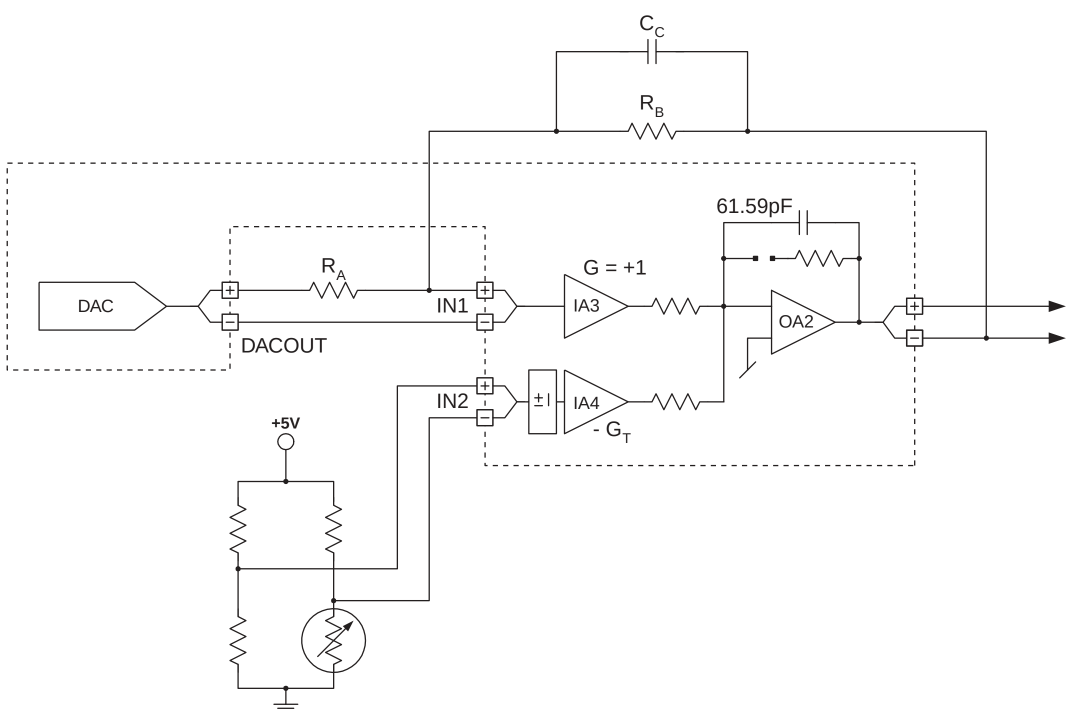

In this circuit, the DC gains are the same as in the high-gain proportional controller, and is given by Equations 5 and 6. The frequency of the dominant pole, and the –3dB point on the controller's frequency response are given by

$$F_{-3dB} = \frac{1}{2\pi R_B C_C} \tag{7}$$

Because of the feedback architecture used in this circuit, the gain from thermistor to output will approach unity, as frequency is increased, as opposed to approaching zero. In many cases, this residual gain will not seriously affect the system stability, and may even improve the system's settling time.

5.3 Integral Control

If high gain can be a good thing, what about infinite gain? A controller with infinite DC gain can potentially provide zero DC system error. A controller based on an integrator (Figure 23a) conveniently provides infinite gain at DC, which gradually attenuates as frequency increases (Figure 23b). Because the 'corner frequency' of an integrator occurs at 0Hz, the responsiveness of an integrator can be characterized either by a gain in front a unity integrator (one in which a 1V input results in a 1V/s output ramp) or by specifying a unity gain frequency in Hz.

Figure 23. Integral Controller (a) and Frequency/Phase Characteristics (b)

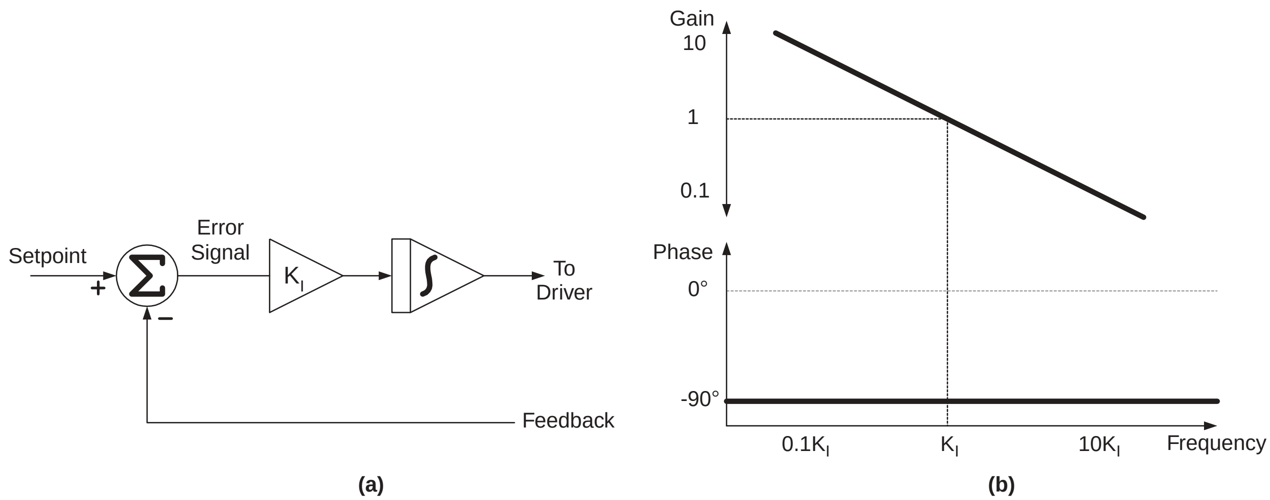

While there are many ways to implement an integrator, one which can be readily achieved with an ispPAC20 is shown in Figure 24. The integration function is achieved by simply opening the resistive feedback path of the OA. The integration feedback is sent back to the input at which the thermistor is sensed. While this does not affect the DC performance of the circuit, it allows the thermistor-output gain to drop below unity as frequency increases, providing a more ideal integrator.

Figure 24. ispPAC20 Integral Controller

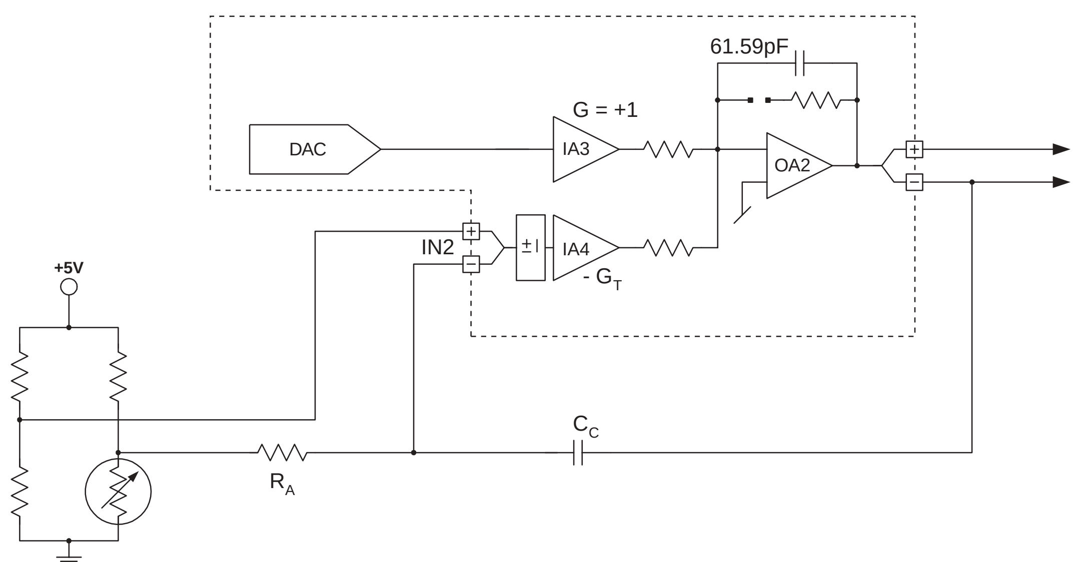

In this circuit, the gain is inversely proportional to frequency, reaching a maximum at DC. While in an ideal integrator, DC gain would be infinite, the DC gain of this circuit is limited by the open loop gain of the PACBlock (~80dB). Assuming that R_A is much larger than the thermistor bridge resistance, the unity gain frequency can be approximated by

$$F_{0dB} = \frac{1}{2\pi R_A C_C} \tag{8}$$

Alternatively, the response of the integrator can be viewed in terms of integral gain (K_I). Note that because this function is an integrator, its response is in units of seconds^{-1}.

$$K_I = \frac{1}{R_A C_C}\ (s^{-1}) \tag{9}$$

While the steady-state error of an integral controller will always be zero, regardless of the integration time constant, the major issue is that of whether the controller will be stable. Compensation for an integral controller is much the same for a dominant-pole controller, keep increasing the time constant until the system becomes stable.

5.4 Proportion-Integral (PI) Control

While slowing down a DP or integral controller by increasing its time constant will tend to stabilize the control system, it also has the often undesirable side effect of slowing down the system's response and settling times. This behavior is a direct result of attenuating the controller's high-frequency response. By sacrificing high-frequency gain to obtain stability, we have traded off the controller's ability to react quickly. In both the DP and integral controllers, the gain is inversely proportional to frequency. This degree of attenuation may not be necessary to maintain stability, but the variables available for adjusting the DP and integral controllers do not allow for setting a minimum gain.

The proportional-integral (PI) control function is obtained by summing a proportional gain and an integrator (Figure 25a), and provides a frequency response that reflects a combination of these two functions (Figure 25b). In a PI controller, the integrator provides infinite DC gain, resulting in zero error, while the proportional amplifier provides a minimum gain at high frequencies. The value of the proportional gain is that it allows the system to settle at a faster rate than it would with an equivalent integral controller.

Figure 25. Proportional-Integral (PI) Controller (a) and Frequency/Phase Characteristics (b)

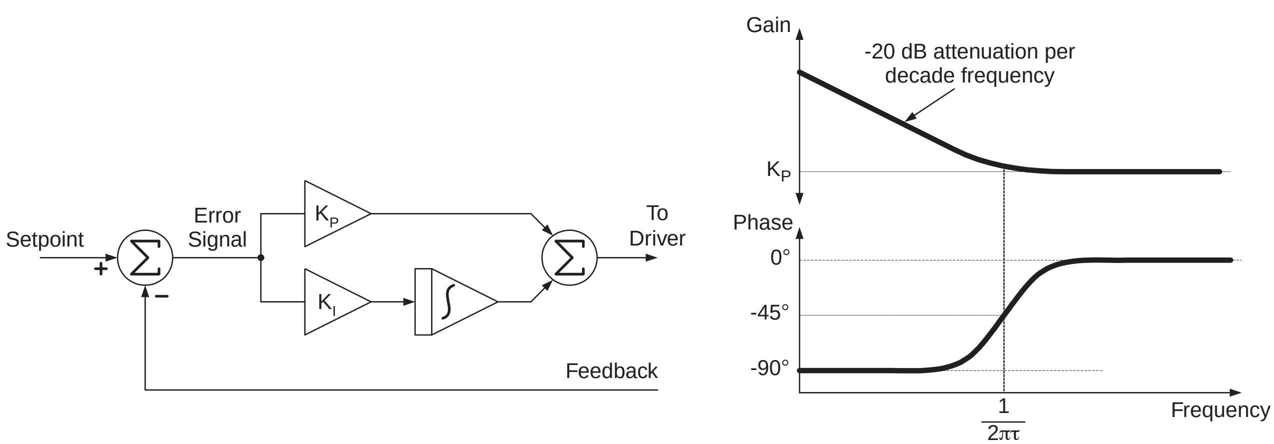

(a) (b)

Figure 26 shows an ispPAC20 implementation of a PI controller. In this circuit, the proportional gain (K_P) is given by

$$K_P = \frac{R_B}{R_A} + 1 \tag{10}$$

and the integral gain (K_I)is given by

$$K_I = \frac{1}{R_A C_C} \ (s^{-1}) \qquad (11)$$

keeping in mind that it cannot exceed the open loop DC gain of the ispPAC20's amplifier (≈80dB). The corner frequency (F_C) below which the integrator starts dominating the response is given by

$$F_C = \frac{1}{2\pi R_B C_C} \qquad (12)$$

Figure 26. ispPAC20 Proportional-Integral (PI) Controller

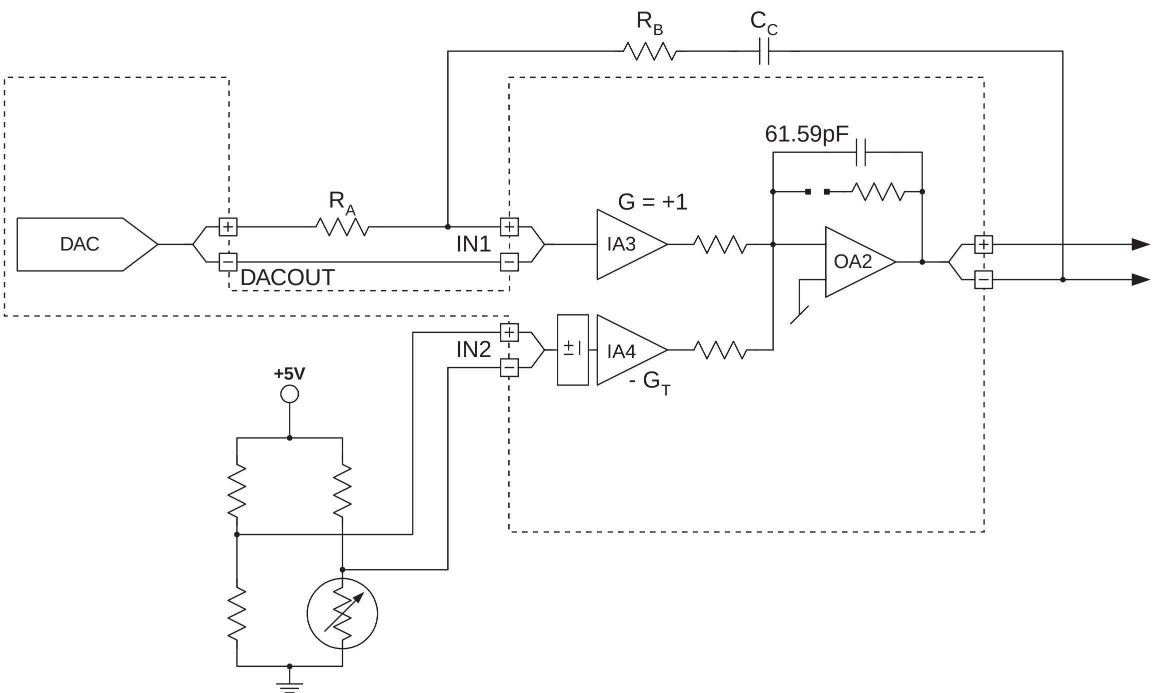

6. Example System

Figure 27 shows a simple, but complete proportional temperature controller based on the ispPAC20 and the circuits presented in this applications note. It is designed to drive TEC modules requiring 2-4V at currents up to about 1A. When used with a type-'J' thermistor, the available temperature setpoint ranges from +15°C to +35°C, with ≈0.1°C resolution. The overall loop gain can be trimmed by varying the ratio of the IA1 and IA2 gains, with IA1=10 and IA2=1 resulting in maximum loop gain.

Figure 27. Minimal ispPAC20 TEC controller system (Proportional)

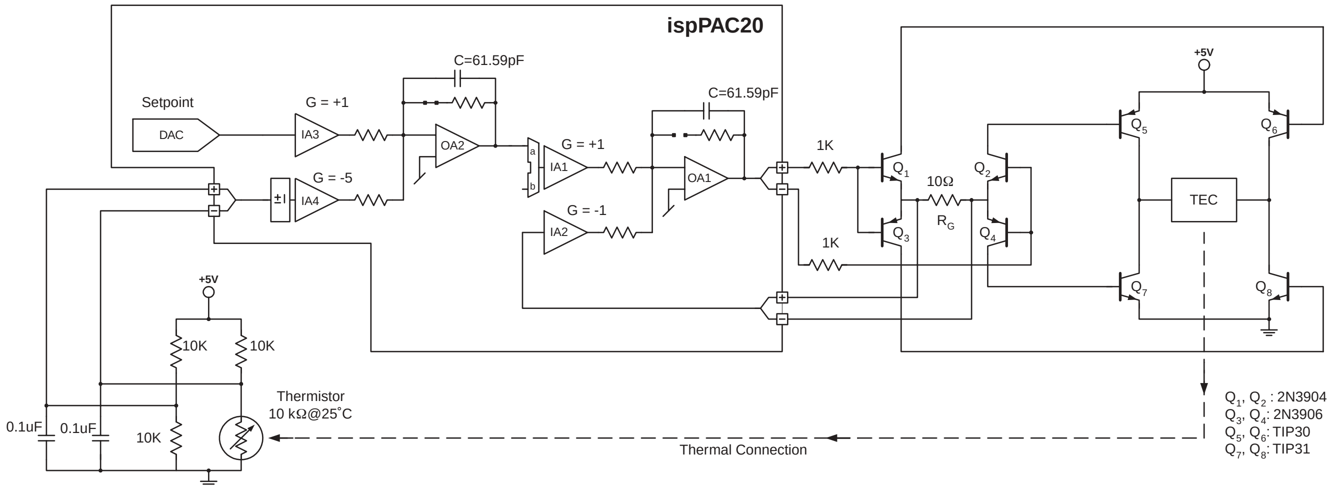

Figures 28 and 29 provide some examples of the behavior of this circuit when controlling a TEC module. In each of these oscilloscope captures, the top trace represents temperature, on a vertical scale of 4°C/div and the bottom trace represents controller error. The horizontal time scale is 10s/div. In Figure 28, the controller gain is set to a low value and one can see minimal overshoot as the controller brings the load to the setpoint temperature, both as the setpoint is raised (Figure 28a) and as it is lowered (Figure 28b).

Figure 28. Response of Proportional Controller, Setpoint Increase (a) and Decrease (b)

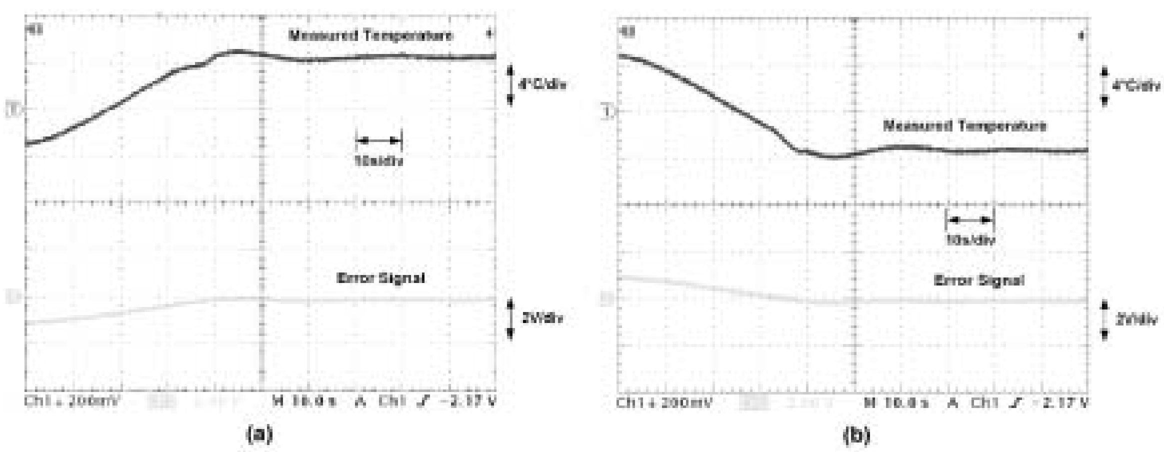

If one increases the gain, one can see a significant increase in overshoot, ripple, and settling time both on a setpoint increase (Figure 29a) and a setpoint decrease (Figure 29b). The ability to rapidly and easily vary control parameters such as setpoint and loop gain is one of the major advantages of using in-system programmable analog circuits. During development, many design alternatives can be explored rapidly. In a production environment, in-system programmable analog circuits allow for individual adjustments of systems on a unit-unit basis if necessary.

Figure 29. Response of Proportional Controller, Increased Gain, Increasing (a) and Decreasing (b) Setpoints

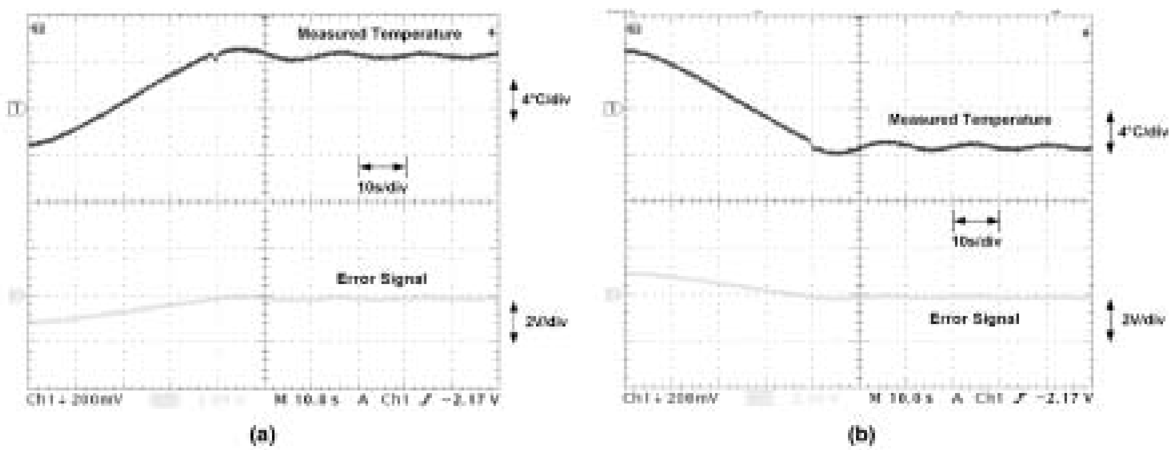

Conclusion

This application note has described how to implement thermoelectric temperature control systems using the Lattice ispPAC20. Circuits for interfacing to thermistor temperature sensors, and circuits for driving the thermoelectric cooler modules were described and analyzed. Several types of linear control strategies (proportional, dominant pole, integral, and proportional-integral) as well as the circuits necessary to implement them were described. Finally, an example of a complete proportional TEC control system based on the ispPAC20 was presented. The flexibility provided by the ispPAC20's programmability allows for simple and efficient implementations of a variety of controllers.

Technical Support Assistance

Hotline: 1-800-LATTICE (Domestic)
1-408-826-6002 (International)
e-mail: ispPACs@latticesemi.com