By Jim Williams, Linear Technology Corp

THERMOELECTRIC-COOLER-BASED TEMPERATURE CONTROLLERS HAVE SOME UNUSUAL REQUIREMENTS. THEY ACCOUNT FOR CIRCUIT- AND THERMAL-DESIGN CHARACTERISTICS TO PROVIDE CLIMATIC PAMPERING FOR TEMPERAMENTAL LASERS.

Controlling the temperature of fiber-optic lasers

Continued demands for increased bandwidth have resulted in the deployment of fiber-optic-based networks. Solid-state lasers drive these fiber-optic lines, which can provide high information density. Highly packed data schemes such as DWDM (dense-wavelength-division multiplexing) involve driving a fiber with multiple lasers to obtain large, multichannel data streams. The narrow channel spacing relies on precise control of laser wavelength to within 0.1 nm. Lasers can provide this precision, but temperature variation influences operation. The intensity of a laser peaks sharply versus wavelength, implying that control within 0.1 nm of laser wavelength is necessary to maintain performance (**Figure 1a**). A typical laser wavelength-versus-temperature plot exhibits a 0.1-nm/°C slope, which means that, although temperature facilitates tuning laser wavelength, the temperature must not vary after the laser wavelength has stabilized (**Figure 1b**). Typically, temperature control of 0.1°C is necessary to maintain laser operation well within 0.1 nm.

A temperature controller for this application must meet some unusual requirements. Most notably, because of ambient-temperature variations and laser-operation uncertainties, the controller must either source or remove heat to maintain control. Peltier-based TECs (thermoelectric coolers), permit this type of control, but the controller must be truly bidirectional. The controller's heat-flow control must not have a dead zone or untoward dynamics in the "hot-to-cold" transition region. Additionally, the temperature controller must be a precision device that can maintain control well inside 0.1°C over time and temperature variations. Laser-based-system packaging is compact, so that size with efficient operation is important to avoid excessive heat dissipation. Finally, the controller must operate from a single low-voltage source, and the controller's power delivery—presumably a switched-mode type—must not corrupt the supply with noise.

TEMPERATURE CONTROLLER MEETS UNUSUAL DEMANDS

The TEC temperature controller in **Figure 2** meets these demands. The circuit includes the input, which includes a DAC and thermistor; a controller stage; and the output stage, which includes the TEC. The LTC1658 DAC, IC_1, and the thermistor form a bridge, and IC_2 amplifies the output. The LTC1923 controller, IC_3, is a PWM that provides appropriately modulated and phased drive to the power-output stage.

The laser is an electrically delicate and expensive load. As such, the controller provides a variety of monitoring, limiting, and overload-protection ca-

Figure 1

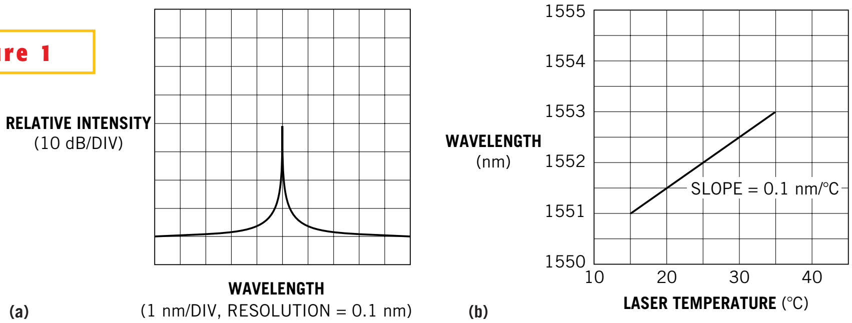

The peak intensity of a laser approaches 40 dB within a 1-nm window (a), and laser wavelength varies by approximately 0.1 nm/°C (b). A typical application requires wavelength stability within 0.1 nm, mandating temperature control.

Figure 2

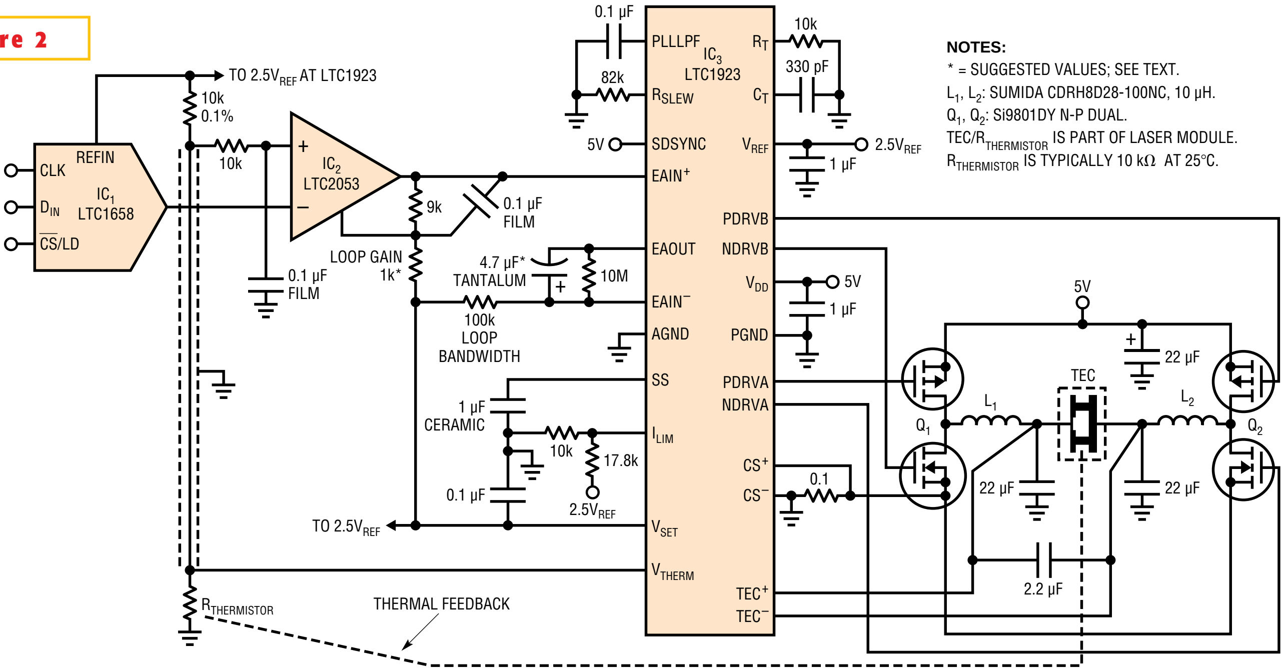

NOTES:
* = SUGGESTED VALUES; SEE TEXT.
L_1, L_2: SUMIDA CDRH8D28-100NC, 10 µH.
Q_1, Q_2: Si9801DY N-P DUAL.
TEC/$R_{THERMISTOR}$ IS PART OF LASER MODULE.
$R_{THERMISTOR}$ IS TYPICALLY 10 kΩ AT 25°C.

A TEC temperature controller includes a DAC, IC$_1$; a thermistor-bridge amplifier, IC$_2$; a switched-mode controller, IC$_3$; and a power output H-bridge. The DAC establishes the temperature setpoint, and you use a gain adjustment and compensation capacitor to optimize loop-gain bandwidth.

Figure 3

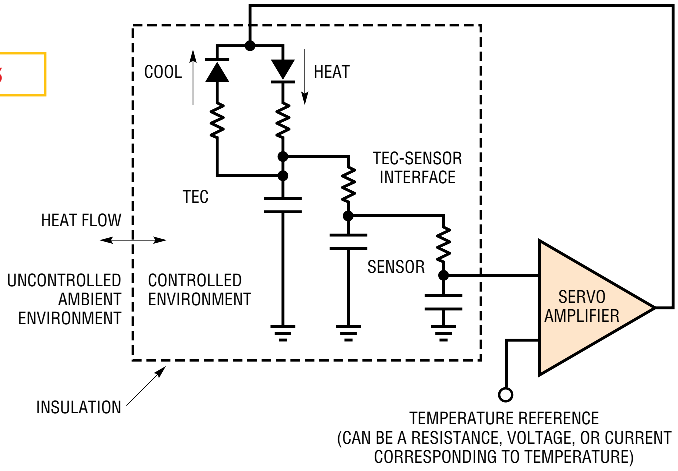

A simplified TEC control-loop model uses resistors and capacitors to represent thermal resistance and capacity, respectively. To avoid instability, the servo amplifier's gain bandwidth must take the lumped delay of the thermal terms into account.

pabilities: soft-start and overcurrent protection, TEC-voltage sense, current sense, and "out-of-bounds" temperature sensing. Aberrant operation results in circuit shutdown, preventing laser-module damage. Two other features promote system-level compatibility. A PLL-based oscillator permits reliable clock synchronization of multiple controller ICs in multilaser systems.

Finally, the switched-mode power delivery to the TEC is efficient, but special design considerations are necessary to ensure that switching-related noise does not reflect back into the host's power supply. The controller includes edge slew limiting, which minimizes switching-related harmonics by slowing down the power stages' transition times. This feature greatly reduces high-frequency harmonic content, which prevents excessive switching-related noise from corrupting the power supply or the laser (**Reference 1**). The switched-mode-power-output stage, an H-bridge type, permits efficient bidirectional drive to the TEC, allowing either heating or cooling of the laser. The thermistor, TEC, and laser, which the manufacturer packages as one module, have tight thermal coupling.

The DAC allows you to adjust the temperature setpoint to any individual laser's optimum operating point, which manufacturers normally specify for each laser. Controller gain and bandwidth adjustments optimize the thermal-loop response for best temperature stability.

THERMAL-LOOP CONSIDERATIONS

The key to high-performance temperature control is matching the controller's gain bandwidth to the thermal feedback path. Theoretically, this matching is a simple matter using conventional servo-feedback techniques. Practically, the long time constants and uncertain delays inherent in thermal systems present a challenge. Both servo systems and oscillators are feedback systems, but one is supposed to oscillate, and the other is not. This unfortunate relationship is very apparent in thermal-control systems.

A model of the thermal-control loop

is simply a network of resistors and capacitors (**Figure 3**). The resistors are equivalent to the thermal resistance, and the capacitors are equivalent to thermal capacity. In **Figure 3**, the TEC, TEC-sensor interface, and the sensor all have RC factors that contribute to a lumped delay in the system's ability to respond. To prevent oscillation, gain-bandwidth limiting is necessary to account for this delay. Because high gain bandwidth is desirable for good control, delays need to be as low as possible. Laser-module purveyors presumably address this issue during manufacturing.

Figure 4

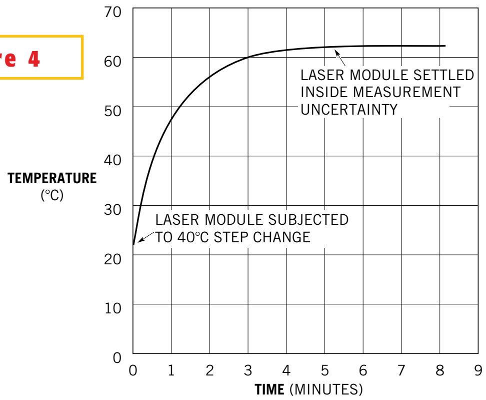

The package thermal resistance and capacity set the ambient-to-sensor lag characteristic for a typical laser module.

The model also includes insulation between the controlled environment and the uncontrolled ambient. The function of insulation is to keep the loss-rate low so that the temperature-control device can keep up with the losses. For any given system, the higher the ratio between the TEC-sensor time constants and the insulation time constants, the better the performance of the control loop (see **sidebar** "Practical considerations in TEC-based control loops").

PRACTICAL CONSIDERATIONS IN TEC-BASED CONTROL LOOPS

A number of practical issues are involved in implementing TEC (thermoelectric-cooler)-based control loops. These issues fall within three loosely defined categories: temperature setpoint, loop compensation, and loop gain.

It is important to differentiate between temperature accuracy and stability requirements. The exact temperature setpoint is unimportant, as long as it is stable. Each individual laser's output maximizes at some temperature (**Figure 1**). You typically increment the temperature setpoint until the laser reaches its peak intensity. At this point, the system requires only temperature-setpoint stability. For this reason, thermistor tolerances are 5% on laser-module data sheets. Thermistor stability over time primarily determines the temperature-setpoint stability over years. Thermistor-time stability is a function of operating temperatures, temperature cycling, moisture contamination, and packaging. The laser modules' relatively mild operating conditions are benign, promoting good long-term stability. You can typically expect thermistor stability comfortably inside 0.1°C over years, assuming that the laser module uses good grade thermistors.

Components that determine the temperature setpoint should have adequate stability over time and temperature. For example, the 10-kΩ, 0.1% resistor in **Figure 2** should have a temperature coefficient of 50 ppm/°C and, preferably, 25 ppm/°C to temperature-setpoint errors approaching 0.1°C over ambient temperature extremes.

An issue related to temperature setpoint is that the servo loop controls the sensor temperature. The laser operates at a somewhat different temperature from the sensor, although laser-temperature stability depends on a stable loop-controlled environment. The assumption is that the laser's dissipation constant remains fixed, which is largely true. This phenomenon also occurs in the sensor's operation. Strictly speaking, the sensor operates at a slightly higher temperature than its nominally isothermal environment. The assumption is that the sensor's dissipation constant remains fixed, which is essentially the case. Because of this fact, the sensor's temperature is stable.

The "dominant-pole" compensation scheme of **Figure 2** takes advantage of the long time constant from ambient into the laser module (**Figure 4**). The loop gain rolls off at a frequency low enough to accommodate the TEC-thermistor lag but high enough to smooth transients arriving from the outside ambient. The relatively high time-constant ratio between the TEC/thermistor and the module insulation, which ranges from less than 1 second to minutes, makes this approach viable. Attempts at improving the loop response with more sophisticated compensation schemes encounter difficulty due to the laser module's thermal-term uncertainties. Thermal terms can vary significantly between laser-module brands, rendering tailored compensation schemes impractical or even deleterious. Note that this restriction still applies, although less severely, even for modules of "identical" manufacture. It is very difficult to maintain tight thermal-term tolerances in production.

The simple dominant-pole-compensation scheme provides good loop response over a range of laser module types. This scheme is the way to go.

Both electrical and thermal gain terms set the loop gain. The most unusual aspect of the loop-gain issue is that the TEC gain differs for heating and cooling modes. Significantly more gain is available in heating mode, accounting for the higher stability in this mode (**Figure 7**). This higher gain means that you should determine loop-gain bandwidth limits in heating mode to avoid unpleasant surprises. (The suggested loop gain and compensation values of **Figure 2** reflect this suggestion.) It is certainly possible to get cute by changing loop-gain bandwidth with the mode, but performance improvement is probably not worth the ruckus. The LTC1923's "heat-cool" status pin beckons alluringly.

A TEC is a heat pump, and the temperature across the TEC determines its efficiency. Gain varies with efficiency, degrading temperature stability as efficiency decreases. Thus, you should provide good coupling from the laser module to a heat sink. Yes, this means you should use that messy white goop. A less obnoxious alternative is to use thermally conductive gaskets, which are nearly as good. The small amount of power involved does not require large heat-sink capability, but adequate thermal flow is necessary. Usually, coupling the module to the circuit's copper ground plane is sufficient, assuming the plane does not already have a thermal bias.

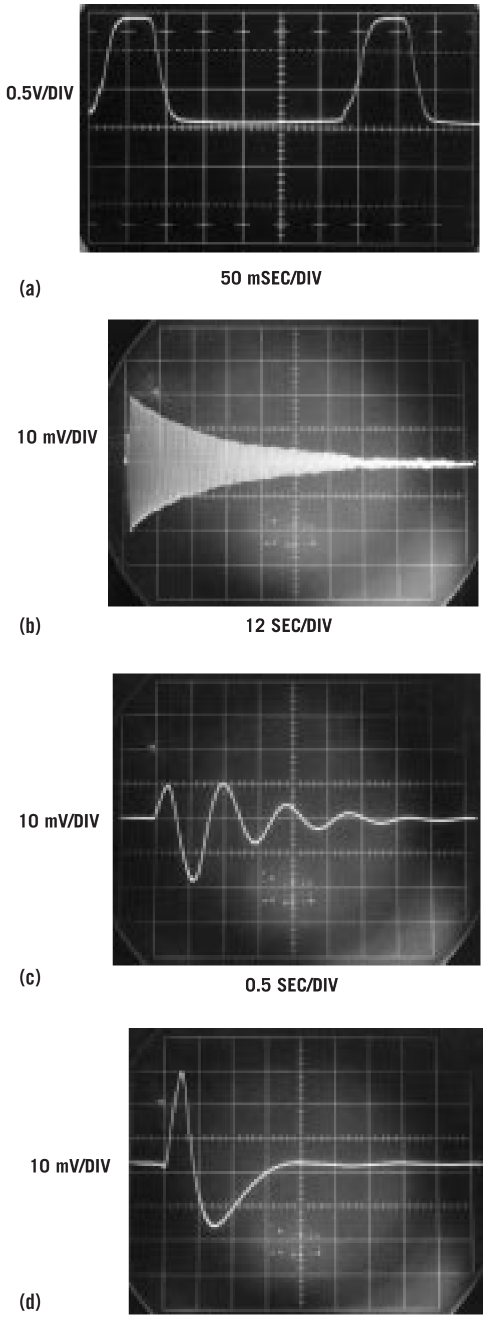

Figure 5

Deliberate excess of loop gain-bandwidth introduces large-signal oscillation, and the duty cycle reveals asymmetric gains for heating and cooling modes (a). When gain-bandwidth is still excessively high, the loop response to a small step in the temperature setpoint results in a damped, ringing response of greater than 2 minutes in duration (b). After reducing the loop gain-bandwidth, the response isn't yet optimal, but settling occurs in 4.5 seconds (c). Gain-bandwidth optimization results in nearly critically damped response with settling in 2 seconds (d).

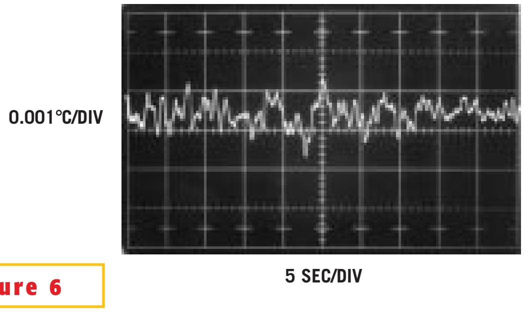

Figure 6

Short-term monitoring in a room environment indicates 0.001°C cooling-mode baseline stability.

OPTIMIZE THE TEMPERATURE-CONTROL LOOP

Temperature-control-loop optimization begins with thermal characterization of the laser module. As mentioned, the ratio between the TEC-sensor and insulation time constants is important. Determination of this information places realistic bounds on achievable controller gain bandwidth. When you subject a typical laser module to a 40°C step change in ambient temperature, the ambient-to-sensor lag, measured in minutes, exhibits a classic first-order response (**Figure 4**). The **figure** plots the laser module's internal temperature, monitored by its thermistor, versus time with no power to the TEC.

You can characterize the TEC sensor's lumped delay by operating the laser module in **Figure 2**'s circuit with the gain set at maximum and with no compensation capacitor. The result is large-signal oscillation due to thermal lag dominating the loop (**Figure 5a**). This **figure** presents a great deal of valuable information. When a circuit "doesn't work" because "it oscillates," whether at millihertz or gigahertz, four burning questions should immediately dominate the pending investigation: What are the oscillation frequency, the amplitude, the duty cycle, and the waveshape? The solution to the problem invariably resides in the answers to these queries. Just stare thoughtfully at the waveform, and the truth will bloom.

In this case, the frequency, which TEC-sensor lag primarily determines, limits how much loop bandwidth you can achieve. The high ratio of this frequency to the laser module's thermal time constant—the lag characteristic in **Figure 4**—means that simple, dominant-pole loop compensation is effective. The saturation-limited waveshape suggests that excessive gain is driving the loop into full cooling and heating states. Finally, the asymmetric duty cycle reflects the TEC's differing thermal efficiency in the cooling and heating modes.

Reducing the controller's gain bandwidth from the extremes of **Figure 5a** produced **Figure 5b**'s display. This waveform results from a small step change in the temperature setpoint of approximately 0.1°C. Gain bandwidth is still excessively high, producing a damped, ringing response that lasts longer than 2 minutes. The loop is just marginally stable. **Figure 5c**'s test conditions are identical to **Figure 5b**'s, but the gain bandwidth is significantly smaller. The response is still not optimal, but settling occurs in approximately 4.5 seconds, or approximately 25 times faster than the previous case. **Figure 5d**'s response, taken at further reduced gain-bandwidth settings, is nearly critically damped and settles cleanly in about 2 seconds. A laser module optimized in this fashion easily attenuates external temperature shifts by a factor of thousands without overshoots or excessive lags.

Further, although substantial thermal differences exist between various laser modules, some generalized guidelines on gain-bandwidth values are possible (see the **sidebar**). A dc gain of 1000 is sufficient for this application's required temperature control, with bandwidth below 1 Hz providing adequate loop stability. **Figure 2**'s suggested gain and bandwidth values reflect these conclusions, although stability testing is mandatory in all cases.

VERIFY TEMPERATURE STABILITY

After you optimize the loop, you can measure the temperature stability by

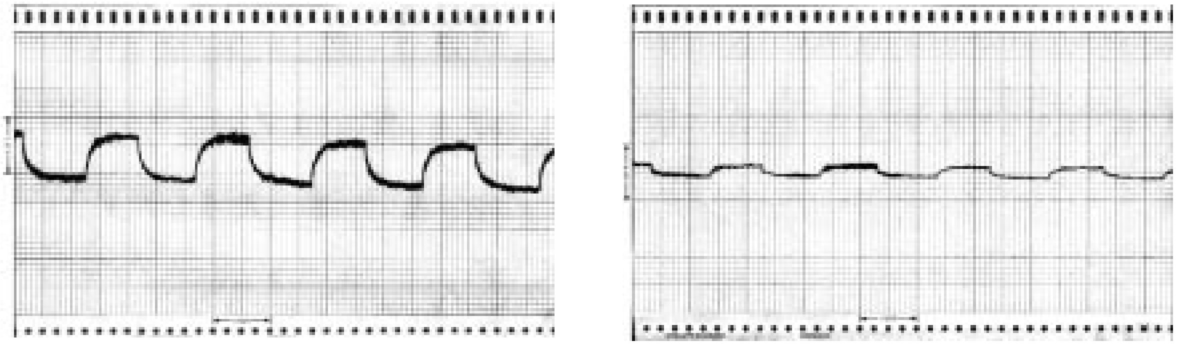

(a) **Note:**
The 0.0025°C baseline tilt over plot length results from varying ambient temperature.

(b) **Note:**
The baseline tilt, which is barely detectable, also shows a four times improvement versus part (a).

Figure 7 A measurement of long-term cooling-mode stability in an environment that steps 20°C above ambient every hour shows a 0.008°C peak-to-peak variation, indicating a thermal gain of 2500 (a). A heating-mode measurement under identical test conditions shows a peak-to-peak variation of 0.002°C, or a fourfold stability improvement due to the TEC's higher heating-mode efficiency, which results in higher thermal gain (b).

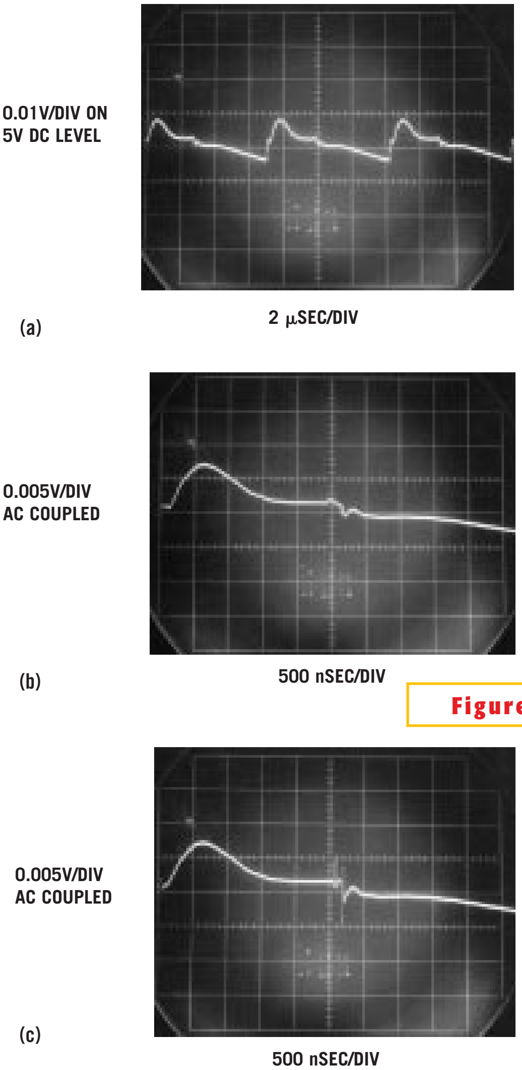

Figure 8 With edge slew-rate limiting in use, the "reflected" noise at the 5V input supply due to switching-regulator operation consists of 12 mV p-p of ripple with much lower high-frequency edge-related harmonics (a). A time and amplitude expansion more clearly shows the residual high-frequency content with slew limiting (b). If you disable the slew limiting, the high-frequency harmonic content rises approximately tenfold (c).

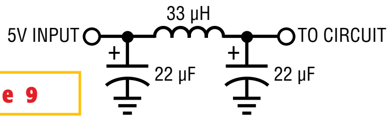

Figure 9 An LC filter reduces reflected ripple to 1 mV and the high-frequency harmonic-noise residue to 500 µV.

monitoring the thermistor-bridge offset with a stable, calibrated differential amplifier. Note that this measurement monitors thermistor stability. The laser's temperature stability is somewhat different due to slight thermal decoupling and variations in laser power dissipation.

Figure 6 records ±1 millidegree baseline stability over 50 sec in the cooling mode. A more stringent test measures longer term stability with significant variations in ambient temperature. **Figure 7a**'s strip-chart recording measures cooling-mode stability against an environment that steps 20°C above ambient every hour over nine hours. (Yes, such archaic devices as strip charts are still useful.) The data shows a variation of 0.008°C, indicating a thermal gain of 2500. (The term "thermal gain" is temperature-control-aficionado jargon for the ratio of ambient-to-controlled temperature variation.) The 0.0025°C baseline tilt over the nine-hour plot length derives from varying ambient temperature. **Figure 7b** uses identical test conditions as **Figure 7a**, except that the controller operates in the heating mode. The TEC's higher heating-mode efficiency furnishes greater thermal gain, resulting in a fourfold stability improvement to about 0.002°C variation. Baseline tilt, just detectable, shows a similar fourfold improvement versus **Figure 7a.**

This level of performance ensures the desired stable-laser characteristics. Thermistor aging characteristics primarily determine temperature stability over years.

The switched-mode power delivery to the TEC provides efficient operation but raises concerns about noise injected back into the host system via the power supply. In particular, the switching edge's high-frequency harmonic content can corrupt the power supply, causing system-level problems. Such "reflected" noise can be troublesome. The LTC1923 avoids these issues by controlling the slew of its switching edges, minimizing high-frequency harmonic content (**Reference 1**). This slowing of switching transients typically reduces efficiency by only 1 to 2%, which is a small penalty for the greatly improved noise performance. **Figure 8a** shows noise and ripple at the 5V supply with slew control in use. A ripple of 12 mV in amplitude is usually not a concern, as opposed to the high-frequency transition-related components, which are much lower in amplitude. **Figure 8b**, a time and amplitude expansion of **Figure 8a**'s display, more clearly studies the high-frequency residue. High-frequency amplitude, measured at center screen, is about 1 mV. A good way to measure the effectiveness of slew limiting is by disabling it. High-frequency content jumps to nearly 10 mV, or almost 10 times worse performance (**Figure 8c**). Leave that slew limiting in there.

Figure 10

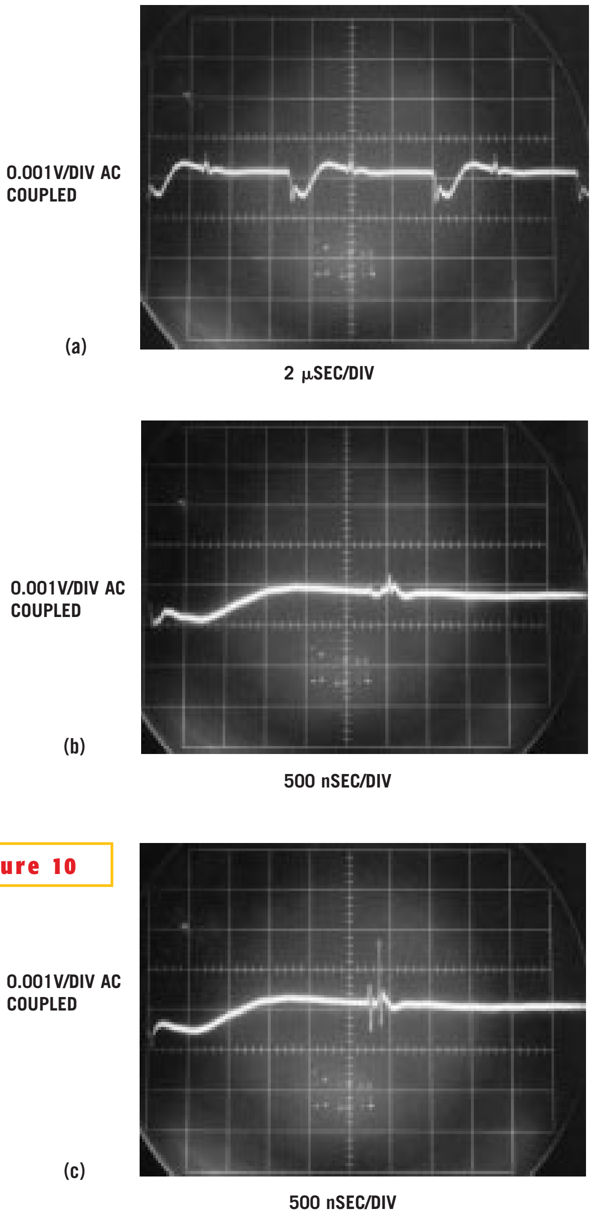

With an LC filter in use, 5V-supply reflected ripple measures 1 mV, and switching-edge-related harmonic content is small due to the slew-limiting action (a). Horizontal expansion shows that the high-frequency-harmonic amplitude with slew limiting is 500 μV, or approximately one-third the amplitude of Figure 9b (b). Without slew limiting, harmonic-content, amplitude rises to 2.2 mV, which is a 4.4-times degradation (c).

This level of noise reduction is suitable for most applications. Some special cases may require even lower reflected noise, and you can use a simple LC filter in these cases (**Figure 9**). Combined with the LTC1923's slew limiting, this filter provides vanishingly small reflected ripple and high-frequency harmonics. With this filter in place, the ripple is only approximately 1 mV, and the high-frequency content is at submillivolt levels (**Figure 10a**). **Figure 10b** expands the time scale to examine the high-frequency remnants. The amplitude is 500 μV, or approximately one-third of **Figure 8b**'s reading. As before, you can measure the effectiveness of slew limiting by disabling it. The result is a 4.4-times increase in high-frequency content to approximately 2.2 mV (**Figure 10c**). So, as before, if you want to achieve the lowest reflected noise, leave that slew limiting in there.□

REFERENCES

1. Williams, Jim, "Switching-regulator design lowers noise to 100 μV," *EDN*, Dec 4, 1997, www.ednmag.com.

2. Williams, Jim, "Thermal techniques in measurement and control circuitry," Linear Technology Corp, Application Note 5, December 1984.

3. Williams, Jim, "Temperature controlling to microdegrees," Massachusetts Institute of Technology Education Research Center, October 1971.

4. Fulton, SP, "The thermal enzyme probe," Thesis, Massachusetts Institute of Technology, 1975.

5. Olson, JV, "A high stability temperature controlled oven," Thesis, Massachusetts Institute of Technology, 1974.

6. Harvey, ME, "Precision temperature controlled water bath," *Review of Scientific Instruments*, pg 39-1, 1968.

7. Williams, Jim, "Designer's guide to temperature control," *EDN*, June 20, 1977.

8. Trolander, Harruff, and Case, "Reproducibility, stability and linearization of thermistor resistance thermometers," ISA, Fifth International Symposium on Temperature, Washington, DC, 1971.

AUTHOR'S BIOGRAPHY

Jim Williams is a staff scientist at Linear Technology Corp (Milpitas, CA, www.linear-tech.com), where he specializes in analog-circuit and instrumentation design. He has served in similar capacities at National Semiconductor, Arthur D Little, and the Instrumentation Laboratory at the Massachusetts Institute of Technology (Cambridge, MA). A former student at Wayne State University (Detroit), Williams enjoys art, collecting antique scientific instruments, and restoring old Tektronix oscilloscopes.