

#20

*PID Control Loops in Thermoelectric
Temperature Controllers*

APPLICATION NOTE

PID Control Loops in Thermoelectric Temperature Controllers

There are many companies that manufacture temperature controllers with various levels of complexity. These instruments use different algorithms and control loops to maintain set temperature. The control loops can range from simple proportional control to proportional-integral control to full PID (proportional-integral-derivative) control. Each type of control loop fills a niche within the realm of temperature control. A temperature controller with full PID capability offers the most flexible control allowing the same controller to bring a very large thermal mass rapidly to set temperature or to maintain very tight temperature stability on a very small thermal mass in an environment with varying temperatures.

This flexibility does come at a price, however. The price is having to determine the appropriate PID values for the thermal system at hand. With many instruments, the determination is done empirically, often after spending a large amount of time trying various combinations of values to see which works best. If this must be done in a production environment where varying loads must be tested at multiple temperatures, the amount of wasted time can grow quickly. The ILX Lightwave LDT-5948 and LDT-5980 offer full PID tuning capability with an auto-tune feature to help remove a large amount of the manual trial and error tuning process and allows it to be automated.

There may be situations however in which the autotune process does not produce a fully optimized temperature control loop for

the thermal load. This Application Note is intended to help describe tuning methods for instruments without autotuning or when autotuning does not provide desirable results. In both of these instances, a better understanding of the PID control loop will allow more efficient determination of the PID constants.

The PID Coefficients

The goal of temperature control is the minimization of temperature error versus the process setpoint. Temperature error e is simply the difference between the measured temperature of a thermal load and its setpoint temperature. In general, the Proportional, Integral, and Derivative components of the complete temperature control effort u_o at time t_o can be described in the following equation:

$$u_o = P \cdot e_o + I \cdot \int_{-\infty}^{\infty} e(t) dt + D \cdot \frac{de}{dt} \quad (1)$$

In this equation, the proportional effort $P \cdot e_o$ is, as its name implies, simply proportional to the temperature error at the moment. A temperature controller with strictly proportional or gain control will have a tendency to oscillate about the temperature setpoint with a fairly long damping time constant if the value of the gain remains constant and set too high. If the gain is set too low, it will take longer than necessary for the temperature to reach the setpoint. In addition, proportional control with constant gain will cause the temperature to settle to a point offset from the actual setpoint leaving a small but nonzero error. This error is generally referred to as “steady state error”.

The integral effort term $I \cdot \int_{-\infty}^{\infty} e(t) dt$ provides a contribution proportional to the amount of time the error has been present. This term allows for offset correction after the proportional effort has decayed away. In cases where the control effort is being limited in some way when there is a large temperature error, for example a TEC temperature controller operating in a current limit mode, the temperature control algorithm can cause the integral effort to “wind-up” or saturate at a value greater than the largest possible effort of the controller. What this means is that when the temperature finally approaches the setpoint, it will overshoot the setpoint by a fairly large amount because the integral effort now has to “unwind”. Controllers should have “Anti-Reset Windup” algorithms incorporated into their designs to prevent this from happening. The 5948 and 5980 are two such controllers. Small values for I allow for minimal overshoot with a long settling tail while large values will allow the load to settle more quickly but at the expense of increased overshoot.

The final term, the derivative or differential effort $D \cdot de/dt$ provides a contribution proportional to the change in error per unit time and prevents control changes larger than the proportional effort when either a setpoint or load change occurs. In essence, the reason for the derivative effort is to provide one large, not necessarily precise, correction immediately after a change in order to start reducing the error as quickly as possible. This term can be used to reduce overshoot as well. Small values of D allow the temperature to change with minimal restrictions while larger values prevent large temperature errors from overwhelming the system response. If the derivative term is nonzero, larger values of the proportional and integral terms may be used to

obtain tighter control.

If the error signal is particularly noisy, the derivative term may induce erroneous “corrections” that are unwanted. Control algorithms should have filtering designed into them to prevent this action. Because the differential effort is only active when the error is changing, a large rapid change in the setpoint can produce a very large “kick” or impulse to the overall control effort. In some cases, this spike in output can be damaging or simply unwanted. In cases like this, the derivative term can be rewritten to incorporate the negative time derivative of the process variable itself instead of the error term. If the process variable is temperature, the new derivative term would become $-D \cdot d/dt (T(t))$. However, if the setpoint fluctuates at all when it should be constant, the differential control effort will produce erroneous results. To prevent these complications from even arising, it may be better to simply nullify the differential term by setting $D = 0$.

Applying a PID Control Loop to a Thermal Load

A TEC controller with one set of PID values will not necessarily control temperature the same way if a different thermal load is connected. Additionally, the thermal characteristics of the load typically necessitate different PID constants if the same load is controlled at different or wide setpoint temperatures. One reason for this is the variation in efficiency a TEC experiences with varying temperature. When a Peltier device is being used to heat a thermal load, the I^2R heating generated by the normal operation of the Peltier junctions increases the efficiency because that heat is pumped into the load in addition to any heat drawn from the heatsink. When cooling a

thermal load, this same Joule heating of the Peltier junctions reduces the module efficiency because this extra heat must be removed in addition to that from the load itself. These differences in efficiency result in different proportional gain values required. For example, a control loop optimized to quickly attain and maintain a setpoint temperature of 45°C may cause the temperature to oscillate if those same PID values are used for a setpoint temperature of 5°C. For this reason, if the temperature setpoint must change more than 10°C, it is recommended that a different set of PID values be used for the second temperature to provide optimal temperature control. While it may not be required for all loads, this guideline can be applicable to most loads where time to temperature is important. If a single set of PID values must be used over several setpoint temperatures, the PIDs generated from the lowest setpoint temperature should be used for all.

Finally, repetitive autotune operations using the same thermal load and temperature setpoint will not usually produce identical PID values. This is because the thermal characteristics of the system do not remain constant from run to run. The temperature of the TEC module, the heatsink, and the environment may and most likely do change from run to run and will present a different thermal system, and hence different PID values.

What determines if a given set of PID values is acceptable for a given application? The answer depends on the test conditions. From the example above, if a test requires frequent changes in the setpoint temperature without constantly updating the PID values, a less than optimal set may be required that results

in a critically damped response in the middle of the temperature range while being underdamped at one extreme and overdamped at the other. Another example may involve the requirement to control the temperature of a device highly sensitive to exceeding some maximum (or minimum) temperature. In a case like this, any overshoot of temperature may cause the device to fail or the sample to be destroyed. This situation would require tweaking the PID values to produce a more overdamped response so that temperature overshoots are minimized. One final example may be the case where many devices must be tested at a specific temperature. These tests would typically fall into a production type of environment. In a case such as this, minimizing the time to temperature would be highly desired so that test throughput can be maximized. Adjusting the PID values to minimize this time may result in some temperature overshoot. As long as the temperature stabilizes to the setpoint quickly, this condition may be acceptable.

System Behavior with Different PID Values

To aid in understanding how different PID values affect a given thermal system, data is presented using an LDT-5948 controlling the internal TEC of a diode laser through an LDM-4984 laser mount. Several different sets of PID values are shown to illustrate how the thermal response changes with varying PID values. Unless otherwise noted, each case begins with a setpoint of 25°C.

In Case 1, the following temperature profile was used:

- 0:00 - Start at temperature setpoint of 25°C
- 0:30 - Change setpoint to 45°C
- 3:00 - Change setpoint to 5°C

5:30 - Change setpoint to 25°C
8:00 - Stop

In the final cases, the temperature profiles were changed to incorporate small temperature perturbations caused by the enabling and disabling of the laser drive current. These cases begin with a temperature setpoint of 25°C. In each case, except the first (where temperature is assumed to be stable), the system is allowed to stabilize for 90 seconds before it is subjected to any additional thermal changes. The temperature profile for these cases was:

00:00 - Start at temperature setpoint of 25°C
00:30 - Enable 100mA laser drive current
02:00 - Disable laser drive current
03:30 - Change setpoint to 45°C
05:00 - Enable laser current
06:30 - Disable laser current
08:00 - Change setpoint to 5°C
09:30 - Enable laser current
11:00 - Disable laser current
12:30 - Change setpoint to 25°C
14:00 - Enable laser current
15:30 - Disable laser current
17:00 - Stop

Figure 1 illustrates the temperature response using the LDT-5948 default PID values of 20.0, 0.800, and 1.000, respectively.

It's obvious that the defaults provide a less than optimal solution for temperature control of this laser. The oscillations cause the TEC to repeatedly swing from positive current limit to negative current limit. This is mostly due to the gain, or proportional, value being set

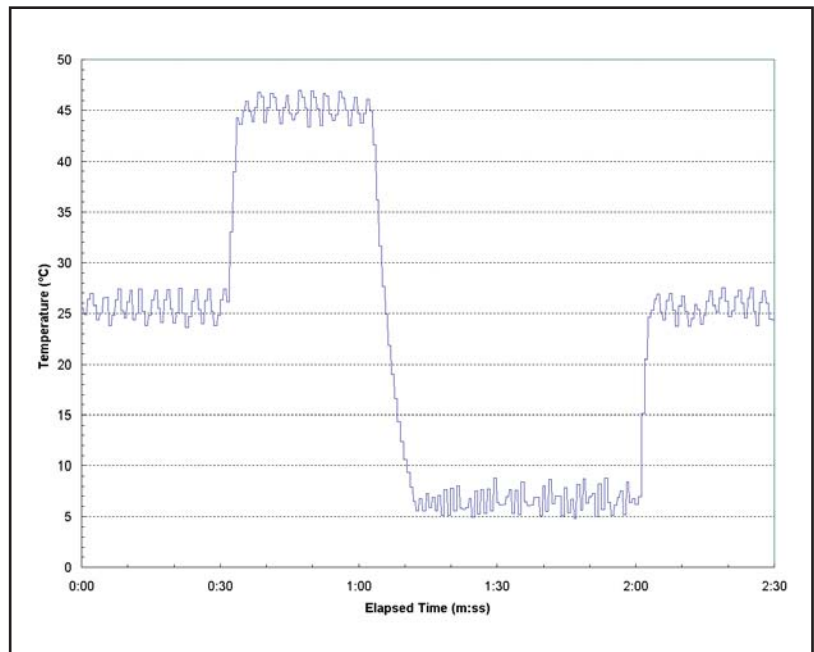


Figure 1. Default PID Values

too high for the very small thermal load being controlled.

In order to temperature control a thermal load where the PID constants are unknown, one may attempt to manually adjust the default PID values to obtain a stable solution. This approach can be very time consuming. A more efficient way is to make use of the instrument's Auto-Tune feature. The instrument will generate a set of PID values to allow rapid closure with minimal overshoot and ringing on the temperature setpoint. The autotune algorithm will typically take anywhere from 15 to 45 minutes to generate the PID values depending on the thermal load. Figure 2 illustrates the result of autotuning with temperature setpoints of 5°C, 25°C, and 45°C.

It may be difficult to determine from Figure 2, but the temperature control performance

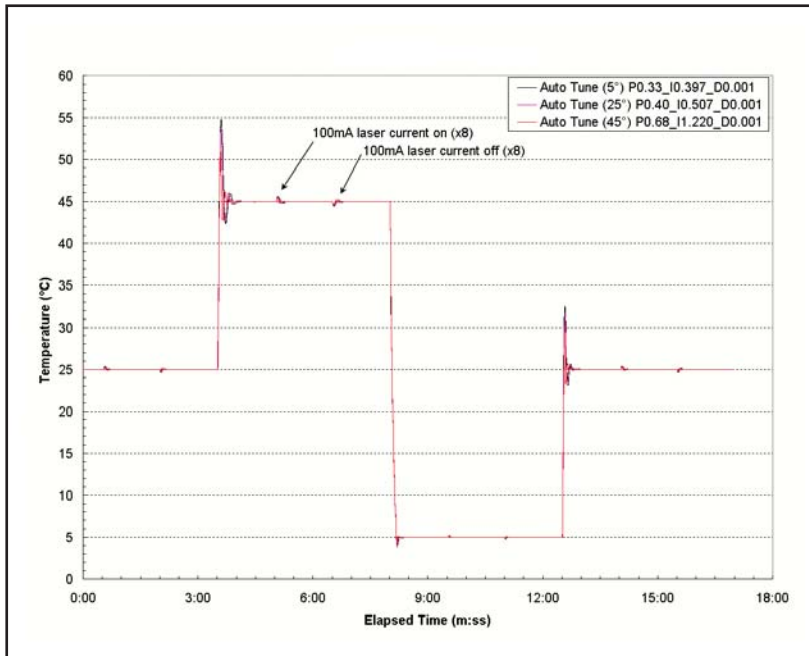


Figure 2. Autotuning Performance

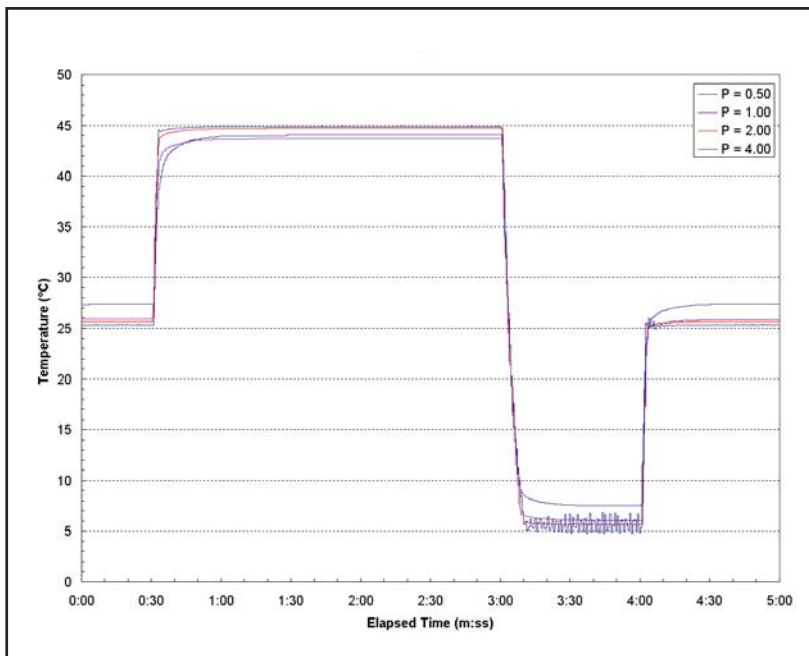


Figure 3. Proportional Variation Only

is affected by the setpoint temperature at which the auto-tune function was run. Figures 6 and 7 show more detail of the temperature settling that occurs at 45°C and 5°C.

As stated earlier, autotuning generates a set of PID values which allow the temperature setpoint to be reached quickly. This implies the load will experience some degree of temperature overshoot and oscillation as it approaches the setpoint. In the case where the load is highly sensitive to thermal spikes or exceeding a maximum (or minimum) temperature, a new set of PID values must be determined manually. To determine PID values manually, it is best to begin by varying only the proportional term with the integral and derivative terms set to zero. The goal is to have as high a gain as possible without causing the output to oscillate.

Figure 3 shows that as the proportional term is increased, the steady-state temperature approaches the setpoint temperature with increasing responsiveness. When the gain is increased too much, oscillation occurs as is evident at the 5°C setpoint with P=4.00.

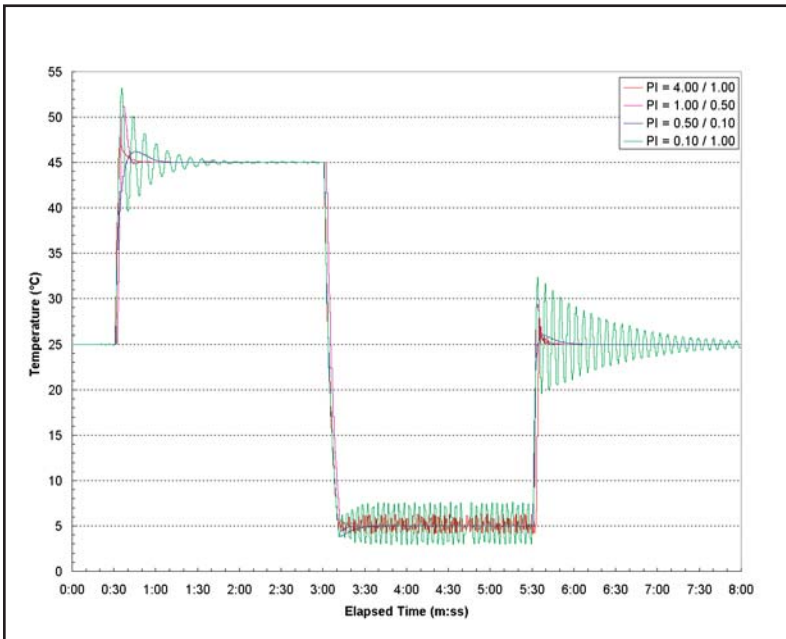


Figure 4. Proportional - Integral Variation Only

Figure 4 illustrates several examples of how the system response changes with different proportional and integral values. In the first case where $P=4.00$ and $I=1.00$, oscillation again occurs at the 5°C setpoint indicating that the proportional gain is set too high. In the other extreme case where $P=0.10$ and $I=1.00$, the gain is set too low which causes oscillations again but this time because the integral term winds up and unwinds in an attempt to control temperature.

Look again at the $P=0.10 / I=1.00$ case. Note the difference in oscillation frequencies and the amounts of dampening that occur between the three temperature setpoints. This behavior indicates three different system responses which would ideally require three different sets of temperature control parameters.

The two intermediate cases show the beginnings of two different control solutions: minimized time to temperature, and minimized overshoot. Adjustments of the control parameters will allow these cases to be optimized.

The examples shown in Figure 5 display the result of varying the proportional and derivative terms. The temperatures have a difficult time reaching their corresponding setpoints since no integral term is present. However, the larger values of D will minimize the temperature error in most cases. This is due to the derivative term providing a larger “kick” at the instant of setpoint change which allows the end temperature to get closer to the setpoint.

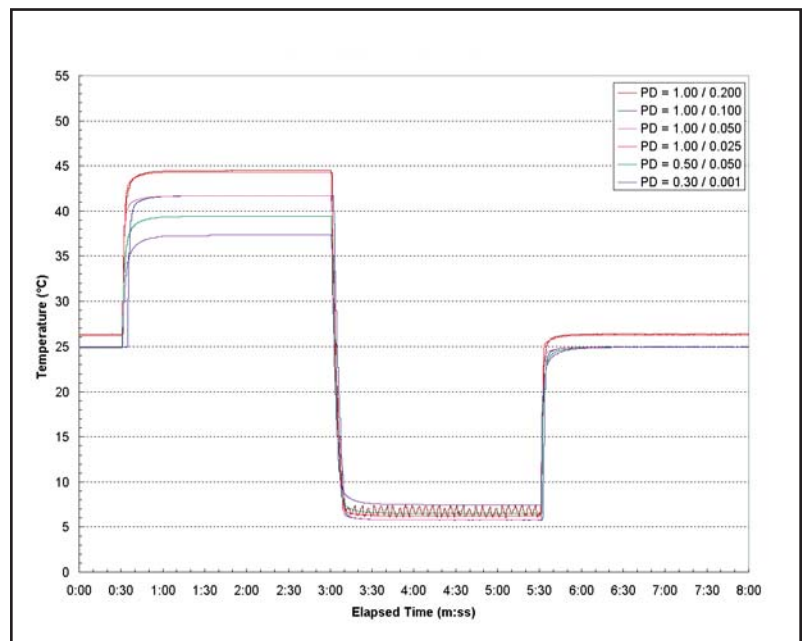


Figure 5. Proportional - Derivative Variation Only

Also, when the derivative term is too large, oscillations can again occur as shown with the 5°C case where $P=1.00$ and $D=0.200$. Variations in system gain with differing temperature are again evident by the fact that one set of control constants which may minimize setpoint error at one temperature fall far from the mark at a different temperature. The case in point is illustrated with the $P=0.50 / D=0.050$ line.

Manual Tuning the PID Values

The proceeding graphs display a limited subset of the variations possible for temperature control when using a controller with a full PID control loop. Because of the six variables (setpoint temperature, thermal load, thermal environment surrounding the load, proportional gain, integral gain, and differential gain) present in a temperature control system, making use of any supplied autotuning algorithm is highly recommended whenever possible. When the coefficients generated from autotuning do not provide an acceptable solution to the temperature control requirements, the coefficients must be generated manually. This process is for the most part trial and error but several guidelines can be given to help accomplish this task:

1. Begin with the integral and differential terms set to zero.
2. Starting with a small value for the proportional term, increase P until the temperature becomes unstable or oscillations just begin. The equilibrium temperature will not match the set point temperature.
 - a. If control at a single temperature is required, small perturbations on the order of $\pm 10\text{-}20\%$ of the setpoint will exercise the control loop to determine oscillatory or unstable behavior.
 - b. If control at multiple temperatures with a single set of PID constants is required, be sure to check the load at all temperatures to identify instabilities and then pick the PID values that provide the most stable output for all temperatures.
3. At the point where oscillations occur from the proportional term only, two different approaches can be taken:
 - a. Measure the period of oscillation (in seconds). This value can be entered for the derivative term. Reduce the proportional gain by 50%.
 - b. Reduce the proportional gain by approximately 30% and then increase the differential gain until small (5%-10%) changes in the setpoint induce underdamped oscillations. At this point, choose a value that provides a critically damped response.
4. Increase the integral term until the system becomes unstable and then back it off by anywhere from 10% to 70% depending on how much overshoot is acceptable. It is better to have as small an integral term as possible.
5. Because all three terms are dependent upon each other, tweaking will undoubtedly be necessary to get the system response that is required. Because of the PID variable dependence, do not change more than one variable at a time when tweaking.

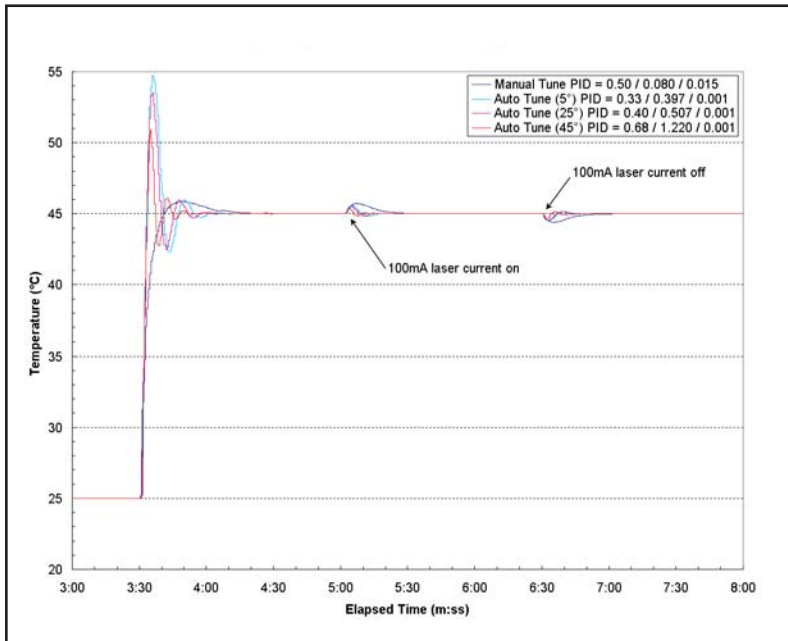


Figure 6. Control Differences at 45°C

Manual Tuning vs. Auto-Tuning

With the goal being to minimize temperature overshoot as temperature is changed between 5°C and 45°C, the PID values were manually determined while following the guidelines in the previous section. The results are compared against results from autotuning at each test temperature in Figures 6 through 8.

The difference in auto-tune optimization temperatures can be seen in Figure 6 where the temperature set-point was suddenly changed from 25°C to 45°C. The overshoot is less and damping period shorter with the 45° auto-tune than with the other auto-tune graphs. The small perturbations created by turning the laser diode on and

off created very similar magnitude responses with all three auto-tune parameter sets. The settling time is slightly shorter with the 45° PID set and slightly longer with the 5° PID set. In all cases, the manual-tuned PID set resulted in an approximate factor of two increase in settling time but with markedly less overshoot when large temperature swings occurred. Very similar results occurred during the transition from 5°C to 25°C.

The transition from 45°C to 5°C showed very little differentiation in overshoot amongst the three auto-tune PID sets. This can be seen in Figure 7.

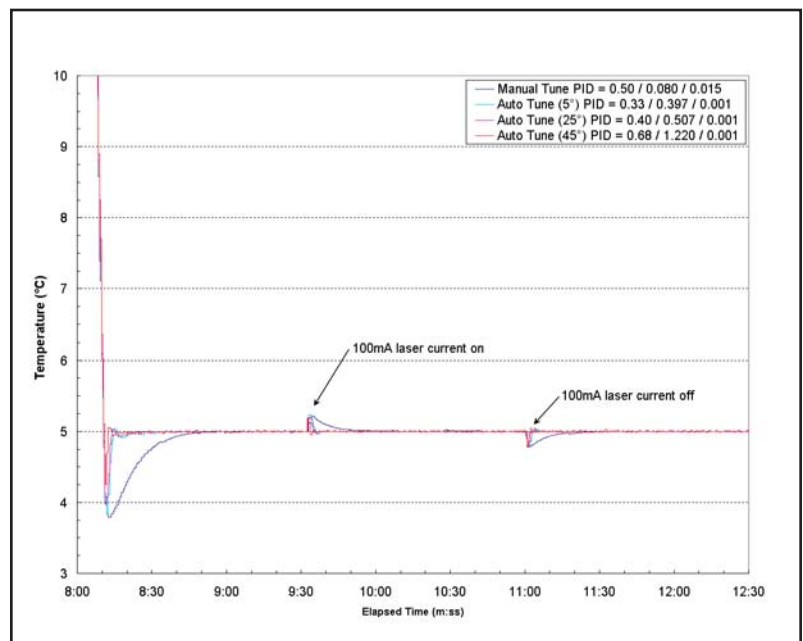


Figure 7. Control Differences at 5°C

For application assistance or additional information on our products or services you can contact us at:

ILX Lightwave Corporation

PO Box 6310, Bozeman, MT 59771-6310

Phone: 406-556-2481 • 800-459-9459 • Fax: 406-586-9405

Email: sales@ilxlightwave.com

To obtain contact information for our international distributors and product repair centers or for fast access to product information, technical support, LabVIEW® drivers, and our comprehensive library of technical and application information, visit our website at:

www.ilxlightwave.com



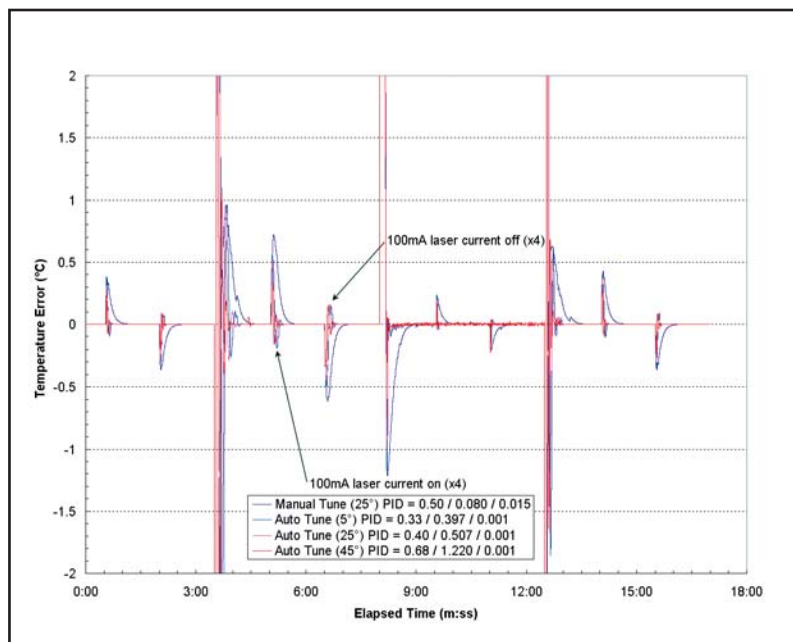


Figure 8. Temperature Error

In fact, the 45° PID set actually has less overshoot than the set optimized at 5°. The main reason the 25° and 45° PID constants are not desirable at 5° comes from the fact that the steady-state temperature is unstable. The manual tune PID set has a longer settling time and comparable overshoot to the longest auto-tune PID case. The manual PID performance could probably be improved if the PID set were optimized at 5° instead of 25°. Figure 8 shows the temperature error for each PID set and temperature setpoint

Conclusion

The procedures and data presented in this Application Note illustrate the following points:

- A temperature controller utilizing a full Proportional/Integral/Derivative control loop is a widely flexible

instrument that can be used with a wide variety of thermal loads.

- Determining the optimal set of PID constants can be a tricky and time-consuming task unless the instrument has an Auto-Tune feature built in.
- Two possible sets of PID constants are possible with a given thermal load and setpoint temperature: minimized time to temperature and minimized temperature overshoot. Depending on the application, one PID solution may work better than the other.
- Even with a carefully determined set of fully optimized PID constants, their use with a different thermal load or simply at a different temperature setpoint may lead to unexpected results ranging from overshoot to oscillation and/or instability.

There are many good sources of information on PID control loops. A good overview is presented in the October 2003 issue of *Control Engineering*.