How Temperature Controllers Work

There's more to thermostat switching than meets the eye.

DR.H.VIRANI

Controlling the temperature is one of the basis industrial process tasks. In its essence, a temperature control system involves sensing the temperature of the plant, and depending on whether the plant temperature is higher or lower than the set point, increasing or decreasing, the heating or cooling accordingly. Of course things are not that straightforward. Some plants are easier to control than others, and some processes require more accurate control than others. Therefore, there is a range of controllers to satisfy these needs, each with their own advantages and disadvantages.

On/Off Controllers

The on/off controller (also known as the bang-bang controller) in its simplest form turns the heat off if temperature increases beyond the set point, and turns heat on if temperature drops beyond the set point. The trouble with this scheme is that if the temperature is hanging around the set point (which is what you want it to do) the controller will be turning on and off very quickly and often, especially if there is a bit of noise in the temperature measurement. The "chatter" very quickly shortens the life of the controller output, and if it is driving large solenoids or actuators of some sort they could easily be destroyed. For this reason nearly all on-off controllers need hysteresis built into the control characteristic (see figure 1).

When a little hysteresis is introduced, the plant is forced to cycle between the two temperatures either side of the set point. Typically the hysteresis is about 0.5°C to 2°C. If you can live with your plant being cycled over a couple of degrees or more, than an on/off controller will solve your problem. However, there's a trap. If the plant is slow to respond, then the cycling can be much larger than the hysteresis. The typical control system has a number of elements that will introduce small time delays (see figure 2).

The cumulative effect is that it might be many seconds or minutes before the controller sees the result of turning the heater on or off. Figure 3 compares the respond of a fast and slow plant with the same on-off controller. Usually if the plant is slow the cycling will be unacceptable and you must look at a more sophisticated controller.

Proportional Controllers

The proportional controller adjusts the heating or cooling in proportion to the temperature error it sees. A typical control characteristic is shown in figure 4, and figure 5 demonstrates most of the properties of the proportional controller. While operating inside the proportional band, the controller applies heat according to the temperature error.

Heat applied = (some constant) X (temperature error) = K(To-T)Watts Where K is the controller gain [units: Watts/°C]. Generally the heat lost from the plant is proportional to the temperature of the plant – *ie* the hotter the plant, the more heat is lost.

Heat lost = (constant)X(plant temperature-room temperature =

(T - Ta)/R

Where Ta is the ambient temperature (room temperature) and R is the thermal resistance to ambient (units: °C/Watt). When the plant temperature is steady the heat applied equals the heat lost. If the loop gain of the system, RK, is very large, then the temperature error will be very small. Typical values for RK are 0.5 to 5 or more for an industrial plant, to greater than 200 for "well behaved " plant, *eg* a calibration bath. Notice too that the larger RK is, the less sensitive the plant is to external influences.

Suppose that the ambient temperature changes by 10°C; the plant temperature will change by:

change in T = 10/(1 + RK)

So high values of RK are desirable.

However, there's another trap here. In any system R and K are not constants. For example, the controller usually has filters to remove high frequency interference, noise, etc. Similarly, the plant takes time to respond to heater power changes – it won't see the power changing in the heater at 120 cycles per second (assuming a 60Hz mains). So both Rand K fall to very low values at high frequency and usually become negative. In fact it is inevitable that if you keep increasing the frequency gain then eventually there will be one frequency where RK = -1. Then we are in trouble because the terms with RK in the denominator become very large and the plant/control system will burst into oscillation at the frequency where RK = -1. This frequency can be anything from a few cycles per minute to many minutes or hours depending on how slow the plant is. The result is that there is a practical limit to the low frequency loop gain RK. This means there will always be a temperature error.

This an obvious fact, really, because to be putting heat into the plant at all, the controller must see that the plant is not at the set point temperature. This temperature error is known as offset. On most commercial controllers there is an adjustment usually called reset or manual reset, which adds or subtracts a couple of degrees to the set point and makes the controller look perfect. It also means that the operator doesn't have to add or subtract a few degrees from the set point every time he or she changes it. The other adjustment on proportional controllers is the controller gain. It is usually expressed

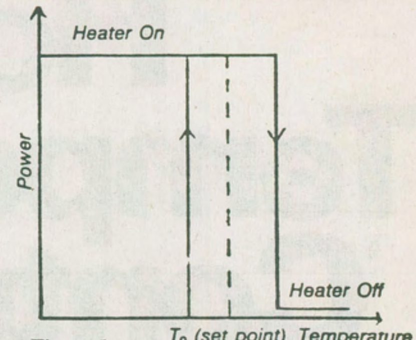

Figure 1

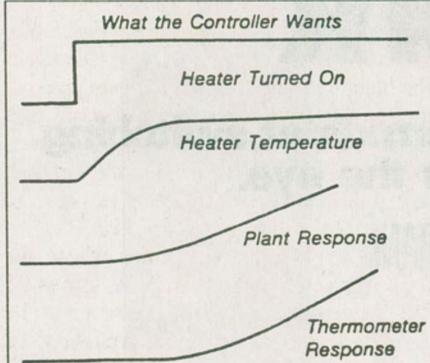

Figure 2

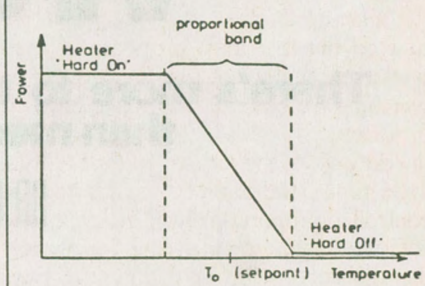

Figure 4

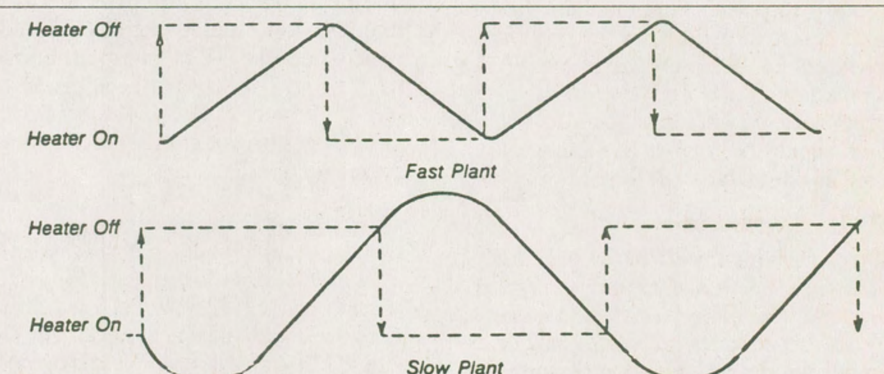

Figure 3

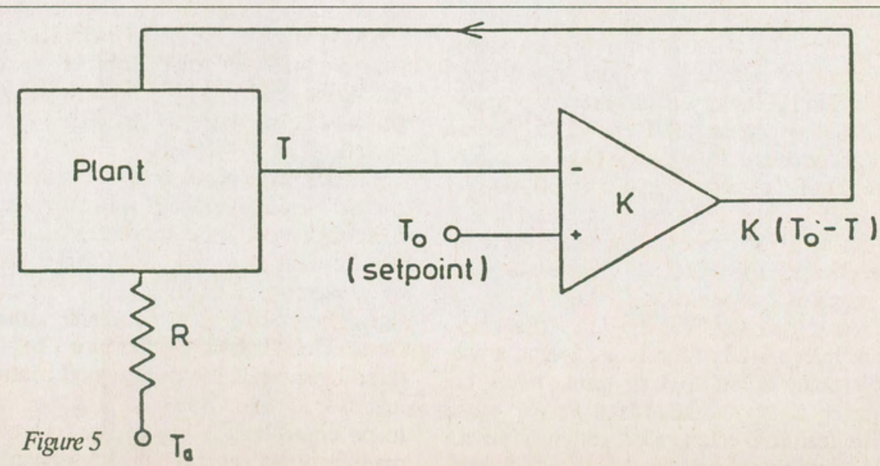

Figure 5

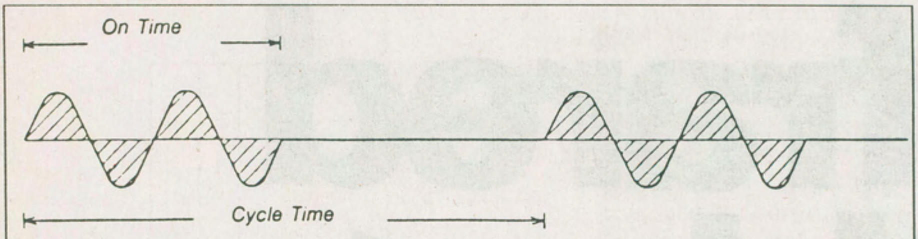

Figure 6

as a percentage of full scale. For a controller with a range of 0 to 1000°C a 5% proportional band is 0.05 x 1000 = 50°C. A 1% proportional band is 10°C. The smaller the proportional band the higher the controller gain.

Having said all that, it should be pointed out that most proportional controllers actually control the power by switching the heater hard on and hard off for short periods, as in figure 6. By varying the duty cycle (the ratio of on-time to total cycle time), the amount of heat can be controlled proportionally. If the cycle time is fast enough, then the plant won't know the difference. Typically the total cycle time will vary between a second for laboratory controllers to minutes for industrial controllers. If your plant is particularly responsive (fast) and the cycle of your controller is long, the cycling will be apparent in the plant response. If the temperature variation proves to be unacceptable a controller with a short cycle time should be sought. Some controllers have an adjustable cycle time.

P.I.D. Controllers

If your plant is particularly slow and difficult to control, it is almost impossible to get the loop gain enough to ensure good immunity to ambient temperature variations. For example, the day and night temperature may differ by 15°C or more. If the loop gain is less than five then the plant temperature will vary at least 3°C each day, and this may be unacceptable. To alleviate this problem some controllers include an integral component (the I in P.I.D.) as well as the usual proportional component (the P in P.I.D.). The integrator constantly adds up the temperature off set (T-To) and applies a correctional signal. Only when the integrator has driven the plant to the set point will the integrator sum stop increasing or decreasing. In this way the integrator component will slowly but surely remove the offset. So long as the offset does not change too fast for the integrator, it will constantly adjust to long term (hours or longer) changes in the ambient temperature. The feature is often called automatic reset. The integrator component then will ensure long term stability and solve the offset problem, but will not help solve the short term stability problem. Ultimately, the only solution is to increase the loop gain somehow. Derivative action (the D in P.I.D. makes this possible). Essentially the derivative component looks at the rate of change of the plant temperature. If it is moving too fast then it says "slow down". This action helps damp out oscillations that may occur when the loop is again very high and the system is marginally stable. It is then possible to operate at slightly higher loop gains.

In some industrial plants where the loop gain might otherwise be as low as 0.5 the derivative component is essential. The big disadvantages of P.I.D. controllers is that they are difficult to tune, particularly if the set point often changes and the plant must settle quickly. On a plant that is inherently slow, tuning could easily take weeks under the eye of an expert. Each of the P, I, and D gains must be adjusted and they all interact with each other. In very difficult plants more sophisticated controllers are required.

Self Tuning controllers

To overcome the P.I.D. tuning problem, a lot of theoretical effort on the part of control theorists has resulted in intelligent "black boxes" called self tuning regulators. They are often P.I.D. controllers with a microprocessor that constantly adjusts the gains to optimize the performance. Others are even more sophisticated. These controllers generally work by constantly twiddling with the gains (in a well defined manner) and assessing the results. Even while twiddling they perform better than any fixed gain P.I.D. controller.

There can be drawbacks with some designs, since when they first start learning to control a particular plant, the twiddling can be excessive and cause major plant disturbances. Nevertheless, they represent the highest performance available with off the shelf controllers. They can only be bettered by special custom designed controllers based on intensive studies of the plant to be controlled. Many of the controller manufacturers now offer such a service. ■