Process Engineers: Take Control!

Cecil L. Smith,
Cecil L. Smith, Inc.

Process control is too important to be left to control engineers. And, dealing with control issues is easier than you might think.

With the ever growing complexity of today's processes and the increasing demands on process performance, the role of process engineers in the design and analysis of process control strategies has to change. This is true both at the initial design stage and in the analysis of control problems in an operating plant.

Two factors are driving the change:

1. Today's emphasis on shortening the design cycle is placing a lot of stress on the traditional approaches to developing the process and instrumentation diagram (P&ID).

2. Corporate rightsizing has reduced, if not eliminated, the cadre of in-house control engineers. Those with the most experience were often the first to qualify for early retirement. Some were immediately hired as consultants, but this pool of expertise becomes diluted with time.

These two factors should force changes that are long overdue.

The integration of process design and process control was deemed to be desirable as long ago as the 1960s. Process engineers and control engineers tend to mix about as well as oil and water, however. They remain mixed as long as you agitate; but as soon as the agitation stops, they quickly separate.

By now, most companies have progressed to some form of digital controls. The available technology includes distributed control systems (DCSs), programmable logic controllers (PLCs), personal-computer-based controls, and microprocessor-based single-loop controllers. The selection of the type of technology and the specific supplier frequently has been a trying exercise akin to a holy war. Most firms have no interest in repeating this endeavor.

Usually, the decisions are based on issues that have little to do with process control. The two features that receive the most attention are the operator interface and the alarm package. (There is no shortage of instant experts on these two topics.) Naturally, the suppliers respond to this attention, and now are introducing features such as three-dimensional graphics. What this has to do with process control is obscure. In reality, these all are system issues, not process issues.

Today, the various approaches and even the available products are largely equivalent. If the job can be done with one, it can be done with any of them. So, why do we need to spend time on the system issues? The time has come to re-emphasize the process issues. This involves two activities:

Design. The key is for the P&ID to truly reflect the nature of the process. Deficiencies in the P&ID cannot be overcome by applying automated tuning, expert systems, or the like.

Evolution. Even with the best of efforts, the initial P&ID will not be perfect. In fact, there are advantages to starting simple and then enhancing the diagram to address the real deficiencies, not the "what if…" and "it would be nice to…" issues. The control strategy also needs to evolve with improved understanding of the process, with new measurement technology, etc.

Understanding the process is truly the key to success in both of these endeavors. Who best understands the process? The process engineers (if not, we are really in trouble). It then follows that the process engineers are in the best position to address the design and evolution of the process controls.

The practice of process control

The reaction of some process engineers to the notion that they should address the control issues may resemble cardiac arrest. This response seems most pronounced among those who, while at university, took the process control course, which usually is a course in mathematics. Process engineers do need an introduction to process control — but not this type of introduction.

Traditional process-control courses place great emphasis on topics such as Laplace transforms. Manuals from control system suppliers occasionally express relationships in the Laplace domain, but they do so for notational convenience. What a practitioner needs to know about Laplace transforms can be taught in 15 minutes. Courses that spend more time than this on Laplace transforms are courses in mathematics.

The keys to process control are the controlled variables, the manipulated variables, the disturbances, and the dependent variables. These are best understood from Figure 1.

The values of the controlled variables and the dependent variables depend upon the values of the manipulated variables and the disturbances. The controlled variables must be maintained at or near a target, although occasionally a range of values for one or more controlled variables is acceptable. Although targets are not provided for the dependent variables, constraints often must be considered. The values of the manipulated variables are at the discretion of the control system, but are possibly subject to constraints. The values of the disturbances are determined by factors other than the control system.

To maintain each controlled variable at its target, the number of manipulated variables must at least equal the number of controlled variables (a "square" system). In most process applications, there are more manipulated variables than controlled variables (a "fat" system), and this creates opportunities for optimization.

Be careful with the terms CV and MV. They could mean controlled variable and manipulated variable. But, alternative terms for controlled variable are process variable, PV, and measured variable, which often is designated as MV. To further confuse things, the controller output sometimes is called the CV. (In this article, we'll limit abbreviations to the familiar P&ID and PID (proportional-integral-derivative) controller.)

Control configurations are dominated by single-loop control, where a controlled variable is maintained at its target by using only one manipulated variable. The control logic is largely the traditional PID control equation. Although the "shrink wrap" has been significantly enhanced through the progression from pneumatic to electronic to digital technology, the PID control equation in today's latest-and-greatest (whatever you deem that to be) is the same as in the large-case pneumatic controllers from the 1940s. The improved shrink wrap offers some opportunities, however, one of which will be discussed subsequently.

Single-loop PID control is adequate for a very high percentage (90% or more) of the controlled variables in most plants. Engineers pay homage to the "keep it simple" principle, but, sometimes, much like politicians honor a "no new taxes" pledge. With all of the fancy features of modern control systems, the temptation to play with the toys is irresistible. The result is "creeping elegance," the most obnoxious effect of which is configuring an excessive number of process alarms.

The nature of the control strategy primarily depends upon the relationships between the controlled variables and the manipulated variables. There are two aspects of these relationships, both of which are determined by the process:

■ Figure 1. The essence of process control.

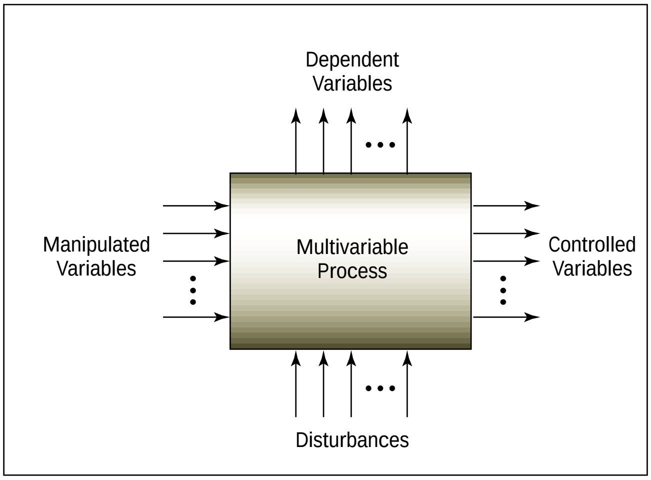

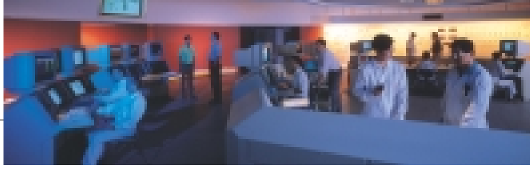

Dynamic. How rapidly does a change in the manipulated variable translate to a change in the controlled variable? Specific attention should be given to dead time. If a change in the manipulated variable has no effect on the controlled variable for some period of time (the dead time), the performance of a traditional PID controller will suffer.

Steady state. After all transients have elapsed, what is the ratio of the change in the controlled variable to the change in the manipulated variable? Control engineers refer to this as the sensitivity or gain of the process.

Traditional process-control courses stress dynamics; the steady-state characteristics, however, usually have the greatest impact.

The three most common process characteristics that lead to control difficulties are:

Nonlinearities (a steady-state characteristic). Envision a plot of the equilibrium values of the controlled variable as a function of a manipulated variable; this is referred to as the process operating line. Tuning difficulties arise when this relationship exhibits a strong departure from linearity. The traditional PID controller is linear; model-predictive control technologies are also linear.

Interaction (a steady-state characteristic). This exists when a controlled variable is influenced by two or more manipulated variables, and, conversely, when a manipulated variable influences two or more controlled variables.

Dead time (a dynamic characteristic). For processes with significant dead time, the performance of the PID controller becomes so poor that alternative model-based technologies such as dead-time compensation or model-predictive control must be considered.

The major nonprocess issue that frequently has a detrimental effect on loop performance is the control valve. Nonidealities such as hysteresis and stiction tend to increase with time, due to the normal wear of mechanical elements. Their impact on loop performance is even larger when the valve is oversized. Too often, plants are saddled with such problems because someone tried to save some money during design and construction. For the same reasons, project managers have led the charge to remove valve positioners. This era hopefully is over, though, with the advent of smart valves.

Smart transmitters and smart valves should be purchased with no discussion. Nonidealities such as hysteresis and stiction do not occur when the final control element is a pump with a variable-speed drive. These are becoming more common, but may never attain preferred status. The measurement devices are our “eyes” to the process; the final control elements are our “handles.” Cutting costs for either invites trouble.

Developing P&IDs

With the traditional approach to process design, the process engineers work up the process flowsheet, from which the control engineers then develop the P&ID. Today’s emphasis on shortening the design cycle is placing a lot of stress on this approach. Even when the process flowsheet is delivered on schedule, the control engineers must quickly gain an understanding of the process, which is no trivial task for a complex process. A late delivery of the flowsheet makes this quite an endeavor.

How do the control engineers go about developing the P&ID? Ask them and you are likely to be told that “you have to understand the process” (or something to this effect). In reality, the control engineers rely very heavily on past experiences, both in what proved to be successful and what did not. This approach works best in those industries like power generation, paper making, and oil refining that repeatedly use the same or very similar processes. But, it also is applied to the one-of-a-kind processes common in specialty chemicals, with far less effective results. Unfortunately, what initially seems to be similar can turn out to be not very similar at all. The likelihood of such mistakes increases as the time allowed for the control engineers to become familiar with the process is reduced.

Less experienced control engineers are being given less time to develop the P&IDs for processes of ever increasing complexity. The limit of these trends is not a pretty picture. The P&IDs for a complex process are developed instantly by a young engineer who just last week was bagging groceries at the supermarket.

As the design cycle is further compressed, a more viable and desirable approach is to have the process engineers designing the process also develop the P&IDs. These engineers certainly understand the process better than anyone else. Developing the P&IDs for a new process involves considerations that usually are more familiar to the process designers than to anyone else, especially at this stage of the project. Will the process designers get it right every time? Certainly not. But, then, the control engineers do not get it right every time either. The question is who can do the job better. The answer is the group with the best understanding of the process — namely, the process engineers.

Now, I am not suggesting that process engineers develop detailed P&IDs, the ones that show practically every component of the process controls. Control engineers generally don’t develop these either. Instead, they work up simplified P&IDs that depict only the key components of the control strategy. Then, the engineering contractor adds the necessary components to create the detailed P&IDs. Process designers certainly can develop the simplified P&IDs.

How do the control engineers go about developing a P&ID? The first step is always “you have to understand the process.” But, hopefully, the process designers already do (if not, inappropriate P&IDs will be the least of our problems). Based upon this understanding, proceed as follows:

• List the controlled variables. Process variable transmitters are required for each of these, although, occasionally, a controlled variable will be computed from other measurements.

• List the manipulated variables. A final control element is required for each. From an instrumentation perspective, the manipulated variable is the signal from the control system to the final control element. But, for control valves and pumps with variable speed drives, process engineers can consider the manipulated variable to be the flow. Where flow measurements are feasible, flow controllers should be installed. From the instrumentation perspective, the manipulated variable is the set point of the flow controller, but, from a process perspective, the manipulated variable is the flow.

• Assess the degree of influence of each of the manipulated variables on each of the controlled variables. For the initial cut, only a qualitative assessment is required. Consider dynamic aspects, especially dead time, as well as steady-state aspects. For each manipulated variable, identify all controlled variables that it significantly affects. Conversely, for each controlled variable, determine all manipulated variables that significantly affect it. A matrix arrangement often is convenient.

• Identify the simple loops. The best situation is when a controlled variable is significantly affected by only one manipulated variable; here, use single-loop PID control (except for some special cases that we'll discuss later). When the controlled variable is affected far more rapidly by one manipulated variable than by any other, single-loop PID control usually proves adequate. For example, dynamic considerations dictate that the bottoms composition for a distillation column must be controlled by manipulating either the bottoms flow or the boilup, not by the overhead flow or the reflux.

• Check for dynamic separation. Interaction is a problem only for loops with approximately the same response speed. If one loop is much faster than another, the fast loop must be tuned first, and then the slow loop can be tuned. The slow loop may not function when the fast loop is in manual, but this is best addressed by improving the reliability of the fast loop.

• Modify the menu of manipulated variables. You can introduce ratios, sums, differences, function generators, etc., to replace one manipulated variable by another. For example, the stack oxygen from a combustion process depends upon the fuel flow and the air flow. The stack oxygen, however, really is a function of the air-to-fuel (or fuel-to-air) ratio. Instead of considering the manipulated variables to be the air flow and the fuel flow, consider the manipulated variables to be the fuel flow and the air-to-fuel ratio (or the air flow and the fuel-to-air ratio). Applying this approach to a process requires substantial insight into the behavior of the process. There is no methodology for determining where ratios, sums, etc., would be beneficial. But, creativity here can be very rewarding.

• Address the multivariable issues. For the remaining controlled variables, the potential for serious interaction exists; this is not necessarily the case, however. Multivariable issues can be approached in two ways:

Process understanding. If you understand the process well enough, you should be able to sort out the interaction issues. Most control engineers take this approach, but their results are not always perfect, which reflects deficiencies in their understanding of the process.

Quantitative measures of interaction (specifically the relative gain). Especially when models provide the basis for the process design, obtaining quantitative values for the degree of interaction is relatively easy. This technology cannot be applied blindly, however. The trick is to understand what these numbers are telling us about the process and our ability to control it.

It is not necessary to eliminate all interaction; it only is necessary to reduce the degree of interaction to the point where the loops will deliver satisfactory performance. Introducing ratios, sums, function generators, etc., also can prove beneficial in this endeavor.

Model-predictive controllers have been suggested as the ultimate solution to such problems. Unfortunately, their implementation requires considerable effort by specialists knowledgeable about the technology. Comparable expertise also is required to keep them running. As originally proposed, model-predictive control was an enabling technology that permitted the benefits of plant optimization efforts to be realized. There certainly are applications where model-predictive control is the preferred (and possibly only) solution, but always identify the benefits (in economic terms) before proposing a model-predictive controller.

• Check for departures from linearity. A change in the manipulated variable will prompt a change in the controlled variable, the ratio being the sensitivity of the process. Will a given change in the manipulated variable lead to a small change under certain conditions, but to a large change under other conditions? This will lead to tuning problems in PID loops, and to a degradation in the performance of a model-predictive controller. A change of 2:1 in the process sensitivity will degrade the performance of the controls, but usually not to a sufficient degree to justify expending effort to address it.

Control techniques are available to address this problem, but each has its limitations:

Cascade control. You must identify a dependent variable that has a linear effect on the controlled variable. Such a dependent variable may not exist.

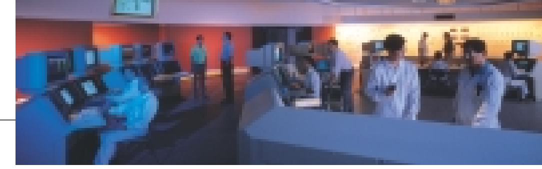

Adaptive control. An understanding of the process is necessary to develop the relationship upon which the adaptive controller is based. Scheduled tuning is a simplistic form of adaptive control, but it also requires some process understanding to be formulated properly.

Self-tuning controllers. These potentially can respond to slow drifts in process sensitivities, due to equipment wear, fouling of heat-transfer surfaces, etc. They, however, cannot respond to rapid changes in process sensitivities.

The best approach is to find a manipulated variable to which the controlled variable is linearly related.

• Assess the impact of disturbances. Where a disturbance has a major influence on a controlled variable, the performance of the control loop will suffer. There are two options:

1. Provide control logic that rapidly responds to the disturbance. The following are possibilities:

Cascade control. You must identify a dependent variable that is affected by the disturbance more rapidly than the controlled variable is. Such a dependent variable may not exist.

Feedforward control. The disturbance must be measured and logic incorporated into the controls to take corrective action before the disturbance affects the controlled variable. Measuring the disturbance frequently is the major obstacle.

2. Eliminate the disturbance at its source. This often is the preferred approach. When the raw materials are natural products, however, some variability is inevitable.

• Consider constraints on the dependent variables. There is one situation where the value of a dependent variable becomes of concern, namely, when it exceeds limits known as constraints. The process controls often include logic specifically designed to respond should the value of the dependent variable approach the constraint. For example, consider a pressure reactor. Provided the pressure is below its constraint, the temperature in the reactor would be controlled. But, should the pressure reach the constraint, the controls switch from controlling the temperature to controlling the pressure. Such logic is referred to as override control.

If violation of a constraint is a hazard to people or equipment, it is the responsibility of the safety system to take whatever action is appropriate should the constraint be violated. For the pressure reactor, the pressure constraint specified to the process controls is less than the setting on the pressure relief device. It is the pressure relief device, not the process controls, however, that ultimately is responsible for protecting people and equipment.

Control problems in operating plants

If the P&ID is appropriate, the plant startup team should be able to successfully commission the controls, which includes such activities as controller tuning. When difficulties are encountered, there are several possibilities. One is that the P&ID is not correct — that is, the control strategy does not properly reflect the characteristics of the process. Basically, there are two alternatives for addressing this:

1. The control engineers can become familiar with the process. Even if such individuals are on staff and currently are available, becoming familiar with a complex process will take some time. If an outside expert is retained, it is usually necessary to reveal proprietary technology.

2. The process engineers can become familiar with the principles of process control. This is where the "just-in-time" education that is possible with computer-based technologies has distinct advantages. (Note the use of the term "education." Training is learning to execute a predetermined sequence of steps. You are trained to change the tire on a car. Education involves mastering principles that can be applied to a variety of situations, some of which are not even contemplated at the time. You are educated to solve process control problems.)

As processes become more complex, the latter alternative becomes more attractive. Indeed, today, process engineers generally are in the best position to address these problems.

The customary initial report of a process control problem is a phone or e-mail message that, in effect, states that some part (or perhaps all) of the controls do not work. If the initial report gives any useful information at all, consider yourself to be fortunate.

As with any type of problem, the first step in its solution is to accurately define the problem. It may be with the routine performance of the controls or with how the controls reacted to a specific event. In either case, we need to address two separate, but related, issues:

1. What did the controls actually do?

2. What was expected of the controls?

Presumably "it doesn't work" means that there is a difference between these two.

A good starting point is to understand what the operations personnel expected the control system to do. This brings up a variety of possibilities:

• The controls responded differently than they were designed to do. There is something in the implementation of the controls or associated hardware (measurement device and final control element) that must be corrected.

• The controls responded as they were designed to do, but this response was not appropriate. The P&ID needs a thorough analysis.

• The controls responded in an appropriate manner. Operations personnel need a better understanding of what control actions are appropriate.

In this article, we'll focus on the first two possibilities.

Determining what the controls actually did involves two aspects:

1. What were the process conditions?

2. What was the response of the process controls?

It is essential to determine these two with a high degree of confidence.

Problems with controls can occur when the process is operating at or near normal operating conditions. A visit to the plant usually is sufficient to observe these problems firsthand and to collect additional data.

Problems also can surface, however, during major upsets or other situations when the process operations are far from normal. This is where a good historian is very useful. Especially during periods of severe process upsets, the recollections of operations personnel can be muddled. The historian is unlikely to provide all of the information required. Key variables may not be recorded at all, others may be recorded at such a slow rate that fast events are missed, and some of the data may not be accurate (measurement devices can provide bad information to the historian). The information from the historian generally has to be supplemented by personal recollections. But, the time-stamped data retrieved from the historian at least are indisputable. It usually is enlightening to understand what they mean.

There are five distinct possibilities for the source of the problem with the controls:

1. a problem with the measurement device;

2. a problem with the final control element, especially if it is a control valve;

3. a problem with the controller hardware;

4. an inability of the process to perform as expected; or

5. a mismatch between the characteristics of the controls and those of the process.

Let's examine each of these.

Measurement device. The potential of a problem with the measurement device must be addressed at an early stage in the analysis of the problem. Does the measured value on which the control action is based reflect the current conditions within the process? If the answer to this question is no, then all other issues are irrelevant. "Garbage in, garbage out" certainly applies.

Rarely can such a problem be resolved merely by calibrating the measurement device. It may have malfunctioned, but, with the built-in diagnostics in smart transmitters, this likely would be detected. Another possibility is that the instrument is doing what it is designed to do, but is not telling us what we really want to know. For example, a temperature transmitter is indicating the temperature of a probe that usually is inserted in a thermowell. What we really want to know is the temperature of the process fluid surrounding the thermowell. These are not necessarily the same. Process engineers generally are most adept at understanding such issues.

Smoothing and filtering also must be examined. An excessive degree of filtering or smoothing on input signals impairs the performance of the controls. The usual approach is that if the input is "bouncing" too much, then increase the degree of filtering. Instead of correcting the root problem, it is easier to hide it with heavy filtering.

In practice, very few inputs require any filtering. The first step is to identify all sources of smoothing. There are three possibilities:

1. within the control system — the value of the smoothing coefficient is always readily available;

2. within the measurement device — with smart transmitters, the value of the smoothing coefficient is readily available; with conventional transmitters, it is considerably more difficult to determine; or

3. between the process and the measurement device (such as partially closing the block valve in the liquid line connecting a displacer level transmitter to the process) — instrument technicians do not hesitate to use such tricks. These can be very difficult to find.

In most cases, the next step is to remove all smoothing.

Final control element. When the final control element is a valve, it very well could be the source of the problem. There are a couple of possibilities:

1. The valve is not positioning to the value provided by the control system; or

2. The relationship between valve position and flow through the valve is very complex.

Although no firm data are available, probably around one-third of all control problems are due to some deficiency associated with the control valve.

Changing the valve position affects the flow through the valve, which, in turn, affects the process. Thus, the valve flow characteristics impact the behavior of the controls. In applications where the flow through the valve is high at times, but low at other times (such as utility and batch processes), the valve characteristics can significantly impact the performance of the controls. If a flow measurement is available, the valve characteristics can be determined by stroking the valve and noting the flow at various valve positions. Otherwise, the valve characteristics have to be calculated from the inherent valve characteristics (supplied by the valve manufacturer) and the pressure-drop relationships for the other components of the flow system. The data for such calculations are more readily available to process engineers than anyone else.

Although considerable attention has been directed to the problem, oversizing of control valves continues to be common. Two issues usually arise with an oversized valve:

1. At large valve openings, an oversized valve has little effect on flow, which definitely will cause problems for the controls.

2. The sensitivity of flow to valve position increases as the valve is oversized. Thus, nonidealities within

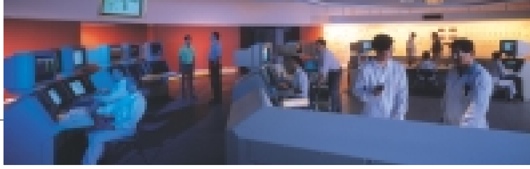

the valve have a larger effect on control system performance when the valve is oversized.

Occasionally, the converse is encountered. As the capacity of the process is increased, flows also must increase. This can lead to some control valves being operated with the bypass partially open at all times. Changing the opening of the bypass valve will affect the performance of the controls.

As the mechanical components wear, actuators exhibit hysteresis. The packing required to prevent leaks around the valve stem resists movement of the valve stem, leading to a behavior called stiction. Both of these usually result in a cycle in both the controller output and the measured variable. When a flow measurement is available, the presence of these phenomena can be easily detected. Are there times when the controller is changing its output but the flow is not changing? If so, then the actuator is not positioning the valve to the value specified by the controls.

Controller hardware. In the days of pneumatic and electronic analog controllers, hardware malfunctions within the controller itself were potential problems that had to be considered. But, today's digital hardware generally is either working perfectly or not working at all. With the self-diagnostic features normally incorporated into computer-based industrial control products, problems with the hardware usually are very obvious.

Both pneumatic and electronic analog equipment are susceptible to significant errors in the calibration of one or more of the controller tuning adjustments. With digital controls, tuning-adjustment calibration errors are simply not possible.

Process. The basic question is "Can the process actually do what the controls are expected to make it do?" With the "our system can do anything you want" sales hype from manufacturers, people sometimes have unrealistic expectations for the process controls. The process ultimately determines what is possible. And, it occasionally does this in some rather subtle ways that only can be appreciated from a thorough understanding of the process itself. Process engineers are clearly in the best position to sort this out.

Controls not matched to process. The performance of the controls is determined by two factors:

1. the control structure, as represented by the P&ID; and

2. tuning of the controls to the process.

The tuning issue must be assessed first.

Tuning. Tuning is the procedure by which the characteristics of the controller are adjusted so that the controller is "in tune" with the process. Tuning techniques have been available since the 1940s. When digital computers were first applied to process control in the 1960s, automated tuning was one of the promises. Yet, despite considerable work and the availability of a number of commercial products, most loops still are tuned with the traditional trial-and-error or "knob twiddling" approach.

The developer of a P&ID does not need to know how to tune a PID controller. But, when analyzing a problem loop in a plant, don't assume that a loop is well-tuned, even if the instrument technicians contend that it is. The analysis of control problems in an operating plant must involve an assessment of how well the controller is tuned.

The first step is to determine the tuning objective for the loop. Figure 2 illustrates two possibilities for the response to a change in the set point: one being the quarter decay ratio (that is, the ratio of the first peak overshoot to the second is 4:1), and the other being critically damped (little or no overshoot). The tuning coefficients are given in Table 1.

Instrument technicians are taught that good performance is a response with a quarter decay ratio. That the quarter-decay-ratio criterion does provide superior performance is illustrated by the responses to a disturbance or load change (tuning coefficients are the same as in the previous example) shown in Figure 3.

A good indication of loop perfor-

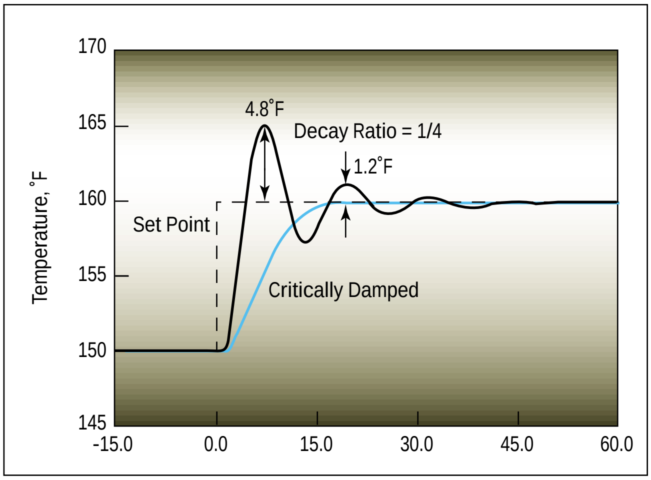

Figure 2. Two tuning approaches.

Table 1. Tuning coefficients for Figure 1.

	Controller Gain, K_c %/%	Reset Time, T_I min
Quarter decay ratio	2.0	5.0
Critically damped	0.5	5.0

mance is the maximum deviation from the set point (the magnitude of the control error at the peak). For our example, the quarter decay ratio gives a deviation of 4.4°F vs. 7.2°F for critically damped.

Production personnel, however, often are not comfortable with the degree of oscillations in responses with a quarter decay ratio. When tuned this way, increasing the sensitivity of any element in the loop by approximately a factor of two results in an unstable loop. For most production personnel, this is a little too close for comfort. When you are involved in plant operations, you do not want your phone to ring at 2:00 in the morning because some loop is unstable. You want dependable controls, and you are willing to sacrifice performance to enhance dependability. After all, you, not the instrument technicians enamored with quarter decay, will get that 2:00 am call.

A widespread malady in controller tuning is for the reset time to be too short, because of the common misconception that fast response is attained through a short reset time. This is not the case at all. The contributions of the three modes are as follows:

Proportional — determines speed of response of the loop.

Integral or reset — forces the loop to line out at its set point (eliminates offset or droop).

Derivative — enhances stability margin, which, in turn, permits the controller gain to be increased to obtain a faster response.

Compared to the controller gain, the reset time has little effect on the speed of response of a loop. To obtain fast response, make the controller gain as high as possible, not the reset time as short as possible.

A major deficiency with using the quarter-decay-ratio performance objective is that, for PI and PID controllers, many combinations of the tuning coefficients will give a response with a quarter decay ratio. In the previous example, a controller gain of 0.5%/% and a reset time of 5.0 min gave a critically damped response. Shortening the reset time to 1.8 min, however, gives a response with a quarter decay ratio. Figure 4 illustrates two responses — both have a quarter decay ratio, but the responses certainly are not equivalent.

The fastest response is the one with the higher controller gain, not with the shorter reset time. The benefits of the faster response are more clearly illustrated by the load responses, as seen in Figure 5 and summarized in Table 2. Note that the faster response (and the one whose maximum deviation from the set point is the lesser of the two) has the

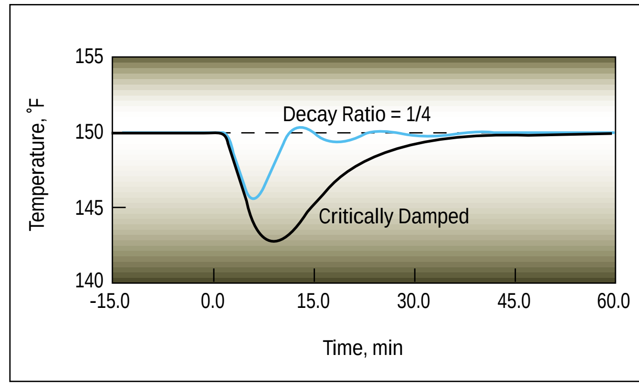

■ Figure 3. Quarter decay ratio provides better response.

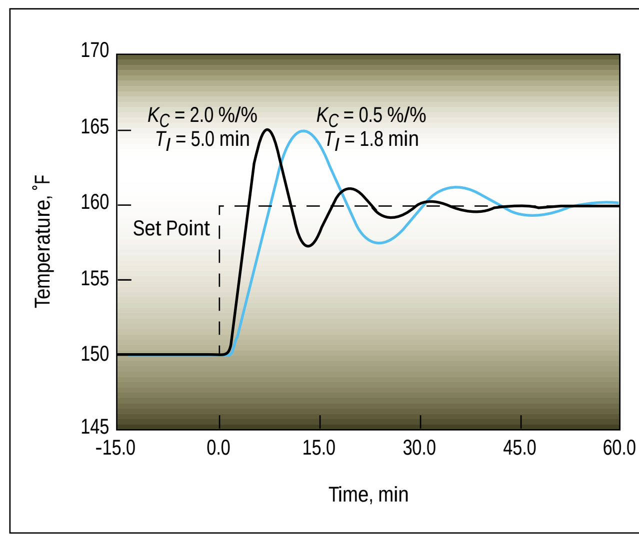

■ Figure 4. Many combinations of coefficients give quarter decay ratio.

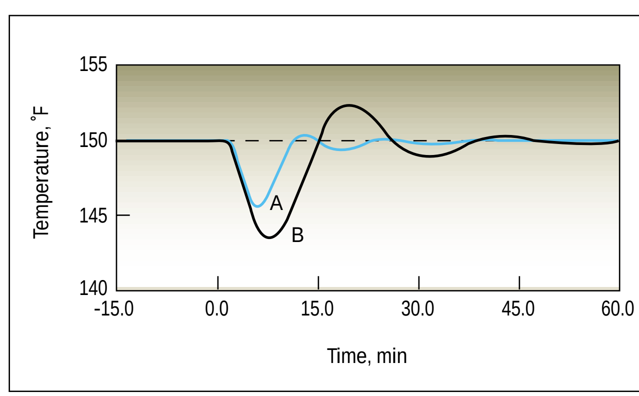

■ Figure 5. Faster response offers benefits.

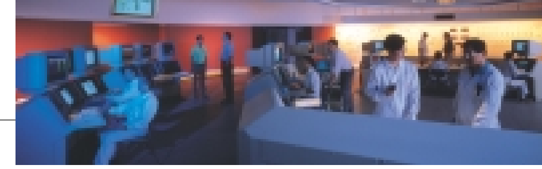

highest controller gain, not the shortest reset time.

Many instrument technicians take pride in being able to tune a loop quickly. One approach to do so is to set the controller gain to some value (how this value is obtained is obscure) and then adjust the reset time to obtain a response with a quarter decay ratio. The technicians consider the loop to be well-tuned. The usual result of this approach, however, is a controller gain that is too low and a reset time that is too short, which results in a slow responding loop. But, instrument technicians don't like to discuss whether a loop whose response is a quarter decay ratio might not be well tuned.

When digital controls are used, there is a way to determine how well a loop is tuned. As noted earlier, a major advantage of digital controls is that there is no error in the tuning parameter adjustments. If the reset time is set to 1.44 min, the reset time is 1.44 min. This statement is rarely true for either pneumatic or electronic controllers.

Suppose a loop has been tuned so that the response exhibits a cycle, although not necessarily a quarter decay ratio. Determine P, period of the response, in min, and T_I, the current value of the reset time, in min. The appropriate value, T_I', for the reset time is given by:

$$T_I' \approx 0.5\, P \tan^{-1} (2\pi\, T_I/P) \qquad (1)$$

If the reset time is less than half of this value, the controller should be retuned. A good approach is to set the reset time to the value computed by Eq. 1, and then adjust the controller gain to obtain the desired performance.

Table 3 presents the values of T_I' computed for the two responses presented previously. This analysis suggests that the loop with a controller gain of 2.0%/% and a reset time of 5.0 min is reasonably tuned, but the loop with a controller gain of 0.5%/% and a reset time of 1.8 min is not.

Because $2\pi\, T_I/P$ is always positive, then

$$0 < \tan^{-1} (2\pi\, T_I/P) < p/2$$

and

$$0 < T_I' < (\pi/2)\, P$$

or, approximately,

$$0 < T_I' < 0.78\, P$$

If you do not like to compute arctangents (although with today's calculators, it is hard to understand why), the following approximation is satisfactory:

$$T_I' \approx (0.78\, P)/[1 + P/(2\pi\, T_I)] \qquad (2)$$

This analysis assumes that the loop should be tuned to respond as fast as possible. Most loops should be so tuned, but, as always, there are exceptions:

Flow loops (specifically those for which the flow meter is immediately upstream or downstream of the control valve). Flow loops are so much faster than all other loops that there rarely is any incentive to make them respond as fast as possible. The typical tuning for a flow loop should be a controller gain of about 0.2%/%, a reset time of about 0.05 min (3 s), and no derivative.

Loops where the flow through the control valve is a disturbance to a downstream process. This most commonly is encountered in level and pressure loops. If the loop is tightly tuned, the disturbances to the downstream process offset the benefits from maintaining the measured variable close to its set point. Such loops have to be tuned with a low controller gain and a long reset time. These loops also are candidates for nonlinear (error squared or error deadband) versions of the PID control equation.

In reality, the following statements apply to controller tuning:

- The value of the reset time solely depends upon the characteristics of the process. This statement also applies to the derivative time.
- Only the value of the controller gain depends upon the performance objective.

Setting the controller gain to some value and then adjusting the reset time to attain the desired performance (the approach often used by instrument technicians) is totally contrary to these statements. The following approach to controller tuning is more consistent with them:

1. With the controller gain set to some value (usually a low value), shorten the reset time until cycling is observed (the cycle does not have to exhibit a quarter decay ratio). Note the period of the cycle and the current value of the reset time.

2. Change the reset time to the value of T_I' computed using the relationships presented previously.

3. Adjust the controller gain so that the response is consistent with the appropriate performance objective (quarter decay ratio, critically damped, or whatever) for the loop.

P&ID deficiencies. Just because a

Table 2. Data for responses shown in Figure 5.

	Controller Gain, K_C	Reset Time, T_I	Maximum Deviation From Set Point
	%/%	min	°F
Response A	2.0	5.0	4.4
Response B	0.5	1.8	6.4

Table 3. Values computed for the two responses.

	Controller Gain, K_C	Reset Time, T_I	Period, P	Appropriate Reset Time, T_I'
	%/%	min	min	min
Response A	2.0	5.0	12.5	7.5
Response B	0.5	1.8	23.0	5.3

loop is specified on a P&ID does not mean that, in fact, it will work. The control structure indicated on the P&ID might not properly reflect the process requirements. The symptom of this is one or more loops that cannot be satisfactorily tuned. The following are the common process problems that lead to untunable loops:

- process nonlinearities;
- excessive dead time;
- improperly nested cascade; and
- interaction between controlled variables.

It is easy to throw technology at such problems, resorting to automatic tuning, expert systems, model-predictive controllers, etc. But, until the problem is understood and the P&ID revised, the loop will not perform satisfactorily.

Process nonlinearities. A nonlinear process exhibits different behavior under different process conditions. Most processes are nonlinear with regard to process throughput. Processes, however, can exhibit nonlinear behavior with regard to other variables as well. Batch processes are particularly notorious in this respect.

In such cases, the controller can be satisfactorily tuned at a point in time. But when process conditions shift, the tuning is no longer satisfactory. The logical starting point is to tune the controller with the process operating at or near normal (or design) conditions. All processes, though, are subjected to major upsets from time to time. How does the controller perform under the process conditions during these upsets? We normally learn the answer through operational experience. This is one source of the "it doesn't work" messages. The first step is to determine the process conditions that led to the problem. The controller can be retuned to these conditions (the "one tuning fits all" approach). This, however, usually leads to some sacrifice in performance at the normal operating conditions. The alternative is to develop a more sophisticated control structure that can cope with the changes in process characteristics.

Dead time. The dynamic characteristic that presents the most difficulty for the PID control equation is dead time. Most processes exhibit dead time, but to a modest extent. This results in some deterioration in the performance of a PID controller, but not enough to justify implementing an alternative approach.

As the dead time rises, the first mode to suffer is derivative. The benefits from derivative decrease as the dead time increases. The next mode to suffer is proportional. As the dead time goes up, the controller gain must be reduced to maintain an acceptable margin of stability. For long dead times, the PID controller essentially will reduce to a pure reset controller. The controller may operate on automatic, but the response is so slow that the controller is practically useless.

For processes with a long dead time, a model-based control technology known as dead-time compensation is available and has been routinely applied to paper machines since the late 1960s. Model-predictive controllers also provide dead-time compensation. There, however, is one prerequisite to the successful application of dead-time compensation: a good value must be available for the process dead time. In paper machines, the dead-time compensator is, in turn, compensated for the effect of changes in machine speed on the dead time.

Improperly nested cascade. In a cascade control configuration, one controller (the outer controller) provides the set point to another controller (the inner one). For this configuration to perform satisfactorily, the dynamics of the inner loop must be faster than the dynamics of the outer loop. The usual desire is that the inner loop be at least five times faster than the outer loop. An inadequate separation of dynamics leads to tuning problems in the outer loop.

When the cascade is a temperature controller providing the set point to a flow controller, such a separation in dynamics almost always is assured. But, when a temperature controller provides the set point to another temperature controller, the separation of dynamics may not be adequate. Furthermore, temperature processes always are interacting stages (each temperature affects the other temperature), which complicates the interpretation of the responses in the inner loop temperature. These control configurations have to be analyzed by someone very familiar with the characteristics of the process.

Interaction. All processes are multivariable in nature. If ten variables are to be controlled, most P&IDs will propose to do this with ten individual loop controllers that are paired with ten final control elements. A single-loop controller works well in a multivariable environment, however, only when the following statements are true:

- The final control element for the loop does not affect the measured variables for the other loops.
- None of the final control elements for the other loops affect the measured variable for this loop.

When these statements are not true, the result is loop interaction. A small degree of interaction can be tolerated, but even modest degrees of interaction will degrade the performance of one or more loops. Sometimes, the interaction can be sufficiently reduced by incorporating ratios, summers, function generators, or other simple elements into the control configuration. For the more severe situations, a model-predictive controller may be the best (or perhaps only) solution.

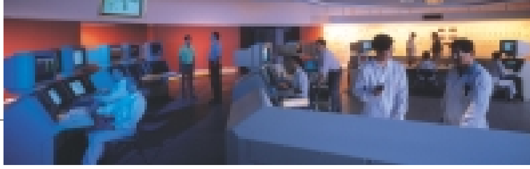

The role for process engineers

Process control issues can be divided into two categories:

Systems issues. These include configuring the hardware, building points (for inputs, outputs, control points, etc.), creating graphic displays for the process operator, and so on. This "grunt work" can be done with very little, if any, understanding of the process. The systems issues are adequately covered by training courses (yes, training, not education) offered by suppliers. Process engineers certainly can do this if they are so inclined, but these tasks should be left to others to free the process engineers' time for the more important aspects of the control system.

Process issues. At the design stage, these center on creating the P&IDs for the process controls. During and after startup, these involve resolving problems with the process controls. Process issues are the important aspects of the control system, and it is crucial that they be done correctly. Both endeavors require a substantial understanding of the process — so, process engineers are in the best position to do them.

Yet, a lot of process engineers do not feel comfortable about addressing the process issues. Many of these engineers probably owe their misgivings to the "process control" course they endured at university — it more than likely emphasized mathematics and gave the wrong impression about the practice of process control.

To address process issues requires an appreciation of the following concepts:

Process characteristics. The focus should be on how steady-state sensitivities and simple process dynamics (integrating vs. nonintegrating, time constants vs. dead times, etc.) impact control loop performance.

PID control equation. Topics include the role of the three modes, the options available in digital systems (especially tracking features), and tuning. Actually, P&IDs can be developed with very little understanding of tuning; tuning must be appreciated, however, by those who resolve control problems in the plant.

Input/output (I/O) elements. This is limited to those aspects of measurement devices and final control elements that affect control system performance. Measurement device selection (thermocouples vs. RTDs vs. pyrometers vs. …) and valve sizing are not required. Some familiarity with filtering and smoothing is needed so that the engineer will not hesitate to remove them.

Split range. This handles situations where, under some conditions, a variable is controlled using one manipulated variable, but, under other conditions, using a different manipulated variable. Most plants have a couple of such applications.

Override control. This technology is employed to address the constraints on production operations. The tracking issues must be understood.

Cascade control. The advent of digital controls has led to an increase in the application of cascade control. Process engineers are in the best position to recognize opportunities for this technology, and, probably, will have to tune most of the temperature-to-temperature cascades.

Ratio control. This simple technology can be understood quite quickly.

Feedforward control. Most process engineers need only an introduction to this topic, although a few should have an in-depth exposure.

Multivariable control. All process engineers must understand how interaction can lead to problems in control systems, and how to recognize when interaction is the source of controller tuning difficulties. Some process engineers must know how to quantitatively assess the degree of interaction and how to develop control configurations that have acceptably low degrees of interaction. Topics such as model-predictive control can be left to the control experts, though.

Let's conclude by putting the proper perspective on mathematics. The PID control equation is a differential equation; the dynamic characteristics of processes are described by differential equations. Engineers do not need to solve either. The control systems solve the PID control equations; various simulation packages are available to solve the differential equations that describe the process. But, in process control, it is necessary to work with differential equations. This is where the Laplace transform enters; as noted previously, differential equations are very conveniently represented in the Laplace domain.

It is understandable why a process engineer's initial reaction would be to attempt to avoid the Laplace transform entirely. The above topics can be presented without using the Laplace transform at all. Laplace expressions, however, actually make the presentation both easier and clearer. In the final analysis, it is better to join them than to fight them.

But, joining them does not require becoming an expert in the subject. Many of the topics (such as partial fraction expansion) taught in college courses never are used in the field. What a process engineer needs to know about Laplace transforms can be conveyed in 15 min. Instead of trying to avoid the subject entirely, it is best to invest the 15 min. Then, you will not be intimidated by the Laplace expressions that occasionally appear in the manuals provided by control system manufacturers. And, it only takes a little knowledge of Laplace transforms to keep the control engineers honest. CEP

C. L. Smith is President of Cecil L. Smith, Inc., Baton Rouge, LA ((225) 761-4392; Fax: (225) 761–4393; E-mail: cecilsmith@home.com), a firm specializing in automation technology. He is the author or co-author of 14 books and over 100 articles. He received BS, MS, and PhD degrees in chemical engineering from Louisiana State University.