An Overview of Proportional plus Integral plus Derivative Control and Suggestions for Its Successful Application and Implementation

By: David Sellers, Senior Engineer, Portland Energy Conservation Inc, Portland, Oregon

ABSTRACT

PID control loops are becoming quite common in the commercial, institutional, and industrial HVAC industry. Proper application and tuning of this control algorithm can bring many efficiency and performance benefits with it. However, improper applications, lack of understanding and poor tuning of these loops are often the root cause behind many commissioning problems. This paper will discuss the background and theory behind PID and why it offers certain advantages from a precision and energy conservation standpoint. It will also look at when the algorithm should and should not be applied. Finally, it would look at how to set up, verify and address control loop tuning issues based on field experience and the work of David W. St. Clair.

INTRODUCTION

Experience over the years seems to indicate that fundamental control system theory and hardware are often misunderstood by HVAC system designers, engineers, and technicians resulting in frustration and the failure of many good designs to fully realize their intended level of performance and efficiency. Consider the following example.

It's early in the morning on a summer day in the Midwest. A young project engineer is performing the final inspection and punch list for a project that implemented extensive renovation and modification of a 45,000 cfm constant volume reheat air handling system serving a hospital emergency and radiology suite. The project had been driven by the need to modify the system to meet the requirements of new radiology equipment installed in the area it served as well as the need to optimize the energy consumption of this energy intensive system. The performance of the discharge temperature controller was critical if the system was going to meet its design goals. Thus the designer had selected and specified a process control grade proportional pneumatic controller believing that this would assure that these requirements would be accurately and reliably met. Even thought the product selected was part of the product line offered by the control contractor for the project, the field staff were not very familiar with its installation, set-up and operation since the controller was infrequently specified due to it's cost relative to the new receiver controllers in the product line.

When the engineer inserted his lab grade mercury thermometer into the system discharge duct to verify the controllers performance he discovered that the system was delivering air at 52.5°F; 4°F below the required set point of 56.5°F. That was 47 tons of unnecessary cooling and 5.7 mbH of unnecessary reheat. He documented the discrepancy as a calibration problem in his punch list and moved on to check other areas. Later that afternoon, when he took the chief facilities engineer to the equipment room to show him the problem, he was further dismayed to discover that the chilled water valve was still in a modulated position even though the day had turned quite hot an humid. Worse yet, the controller was now maintaining a discharge temperature of 60.5°F, even though its set point was unchanged. It looked like that calibrated or not, his sophisticated and expensive controller was not capable of holding a set point.

As a result, an angry letter was written to the control contractor branch manager complaining about the poor quality of both their product and their field support. The branch office responded by sending their best pneumatic control pipe fitter out with a replacement controller. What nobody at the time seemed to realize was that the product was in fact working properly and was probably reasonably well calibrated. Despite everything the project engineer had learned in developing the design for the system modification, despite all his calculations and analysis, despite the combined experience of the branch manager and control fitter, they all had failed to understand the fundamental concept behind a proportional controller; i.e. the output of the controller is proportional to the difference between the set point and the control point.

PROPORTIONAL OFFSET OR ERROR

Stated another way, except for one very specific load condition, there will always be a difference between the control point and the set point for a system that is operating under the control of a proportional controller. This difference between what you want and what you are getting is called the proportional offset or proportional error. How big this difference will be is a function of the gain of the controller. Stated mathematically:

$$Output = K_P * Proportional\ Offset$$

Where:

K_P = controller proportional gain

For many people, an easier way to think of controller gain is to think of it in terms of controller throttling range or controller proportional band. The terms are reciprocals of each other, at least in the general sense. The throttling range of the controller is the change in input, which will cause the controlled device to go through its full stroke. For instance, a pneumatic controller temperature controller with a 10°F throttling range that is controlling a normally closed valve with a 3-15 psi span would require a 10°F temperature change at its input sensor to generate a 12 psi change at its output and make the valve go from fully closed to fully open. A typical controller calibration procedure adjusts the controller so that its controlled device is at mid-stroke when the control point is exactly equal to the set point. So, if the controller we were describing in the preceding sentence was an air handling unit discharge temperature controller which modulated a chilled water valve to maintain a 56.5°F set point and had been calibrated so that the valve was at mid stroke when the controller output was at 9 psi (half way between 3 and 15 psi), then the discharge temperature would need to fall 5°F below set point to fully close the valve and rise 5°F above set point to fully open the valve.

Thus, at low load conditions, the discharge temperature would tend to drop below the required set point until the valve was in a position where the flow through it exactly matched the low load condition. When there was no load, the system would have to allow the discharge temperature to drift down to 51.5°F to completely close the valve. Similarly, at high load conditions, the system would have to allow the discharge temperature to drift up to 61.5°F before the control valve would be driven fully open and allow the design chilled water flow to enter the cooling coil.

It is quite likely that the controller in the case study at the beginning of the paper was a proportional controller with a 10°F throttling range working as designed over a range of load conditions. Early in the day, when the load was low, a stable operating point was achieved with a discharge temperature 4°F below the controller's set point. As the load increased, the chilled water flow at this valve position was inadequate and the discharge temperature started to rise. The rising discharge temperature changed the offset in the system and thus output of the controller. This caused the valve to open until a new point of stability was achieved.

After thinking about all of that for a while, it is a common reaction for less experienced people to suggest that the throttling range be made as small as possible so that the system would have little if any proportional offset, regardless of the load condition. This is in fact one of the parameters that you are trying to achieve when you tune a controller; i.e. you are trying to achieve the smallest possible throttling range, which *will provide stable system performance under all operating conditions*. It is in the second part of that statement wherein the trick and the problem lies.

As you narrow the throttling range of a controller, the

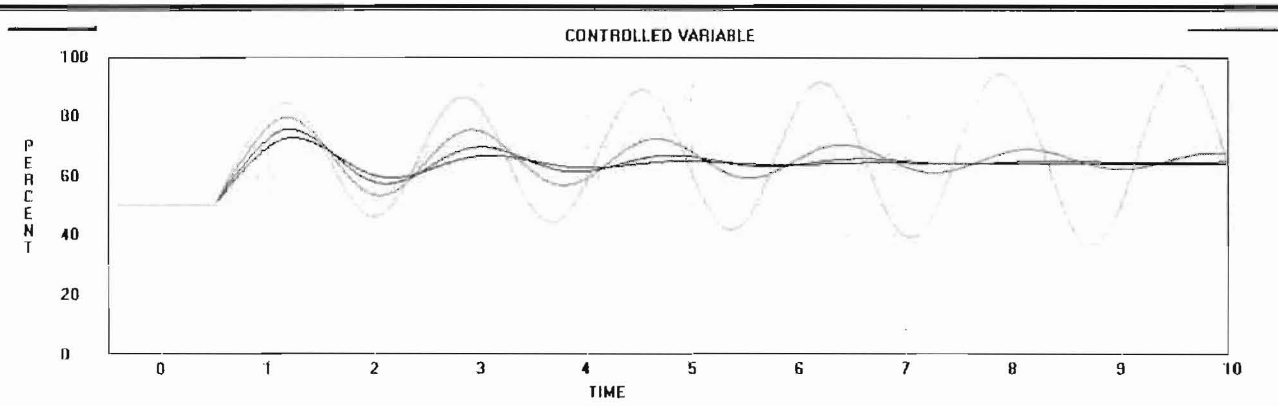

Figure 1 – The Effect of Narrowing the Throttling Range On a Controller

The black line (darkest) represents the response of a well-tuned proportional control system to an upset or change in the system, like a set point change for instance. Notice how there are a few oscillations and then the system finds a new stable operating point at a different valve position. The blue line is the response of this same system after the throttling range has been decreased to the point where it starts to become unstable. The red and the gray lines are the results of further decreases in the throttling range. Note that in the case of the gray line (lightest, which represents the most sensitive system), the magnitude of the swings is increasing.

response of the controller to a change sensed by its input sensor becomes much more pronounced. For instance, a pneumatic temperature controller with a standard 12 psi output range and a 10°F throttling range will adjust it's output 1.2 psi for every 1°F of temperature change at its input. If the controller had a 5°F throttling range, a 1°F temperature change at its input would result in a 2.4 psi change in output. With a 2°F throttling range, the output would change 6 psi for every 1°F temperature change at the input sensor. So, as the throttling range of the controller is decreased, the controller becomes much more sensitive to changes at its input and responds much more strongly to them. This is good as long as the system that the controller is controlling can keep up with the controller's changes. But, as you narrow the throttling range, most real systems will reach a point where time lags, thermal inertia, physical inertia and other factors result in the over-all system response is slower than the controller's response to a change at its input. As a result the system becomes unstable and starts to hunt. (In human terms, the system can't keep up with the demands of the controller and goes crazy.) Figure 1 illustrates what happens as the throttling range is reduced on a system.

To understand this a little better, lets look in detail at what happens with our chilled water valve controller when something causes the temperature to change at its input sensor. We will continue to discuss this as if the controller were a pneumatic unit. The essence of what happens will be the same regardless of whether the controller is a pneumatic instrument, a solid state instrument, or a computer. The first thing that happens is the actual air temperature changes. The change in air temperature causes the temperature of the sensing element to change. There is some time lag associated with that process, but probably not much since sensors are designed to respond fairly rapidly to temperature changes. The controller must now act on this temperature change. In pneumatic controllers, this typically was accomplished via a series of levers, bellows, and other mechanisms that somehow caused air from the air source to be added or removed from the output line. Again, this happened fairly quickly, but some finite amount of time will be required for all the levers, bellows, and other mechanical devices to move to reflect the new operating condition. (See Figure 2[i]).

Now that the controller has made some internal change to try to change its output pressure, the pressure has to actually change before the system will respond. This is a function of several things including the nature of the pressure source, the length of the line to the valve that is being controlled, and the volume of air required to fill the valve actuator enough to start to move the valve stem. Of course, just because the valve linkage starts to move doesn't necessarily mean that the valve plug is going to move. Most mechanisms, especially linkage systems, have some play in them. So, for the valve plug to move and start to actually change the flow through the chilled water coil, the valve linkage must first move enough to take up this play (sometimes termed hysteresis or dead band).

And of course, all of this takes time. Once the play is out of the system and the valve plug actually moves, the flow through the chilled water coil will start to change. But, the higher chilled water flow rate must first cool the tubes and fins of the coil to a slightly lower temperature before the tubes and fins can cool the air stream. Once the air stream starts to cool, there will be some period of time (usually called transport time or dead time) between when the cooler air leaves the face if the coil and when it reaches the sensing element. Once the cooler air reaches the sensing element, it can change the temperature of the sensing element, which will effect the controller output and the cycle repeats. In small systems with fast controllers and rapid fluid streams, the combined impact of all of these time factors will amount to fractions of a second or seconds. On a large system with large final control elements and large distances between the controllers, sensors and coils, these interactions can take fractions of a minute or even minutes.

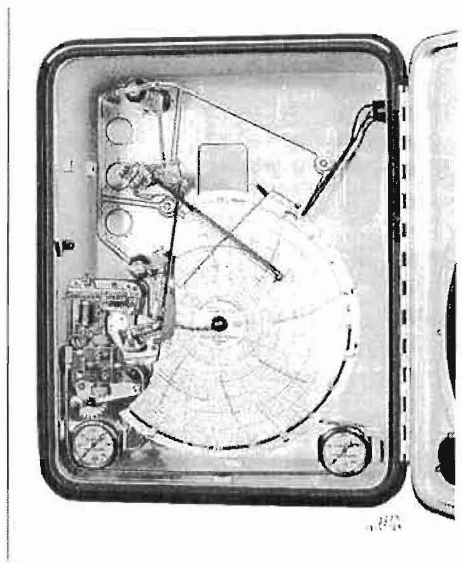

Figure 2 – Mid 1940's State of the Art Pneumatic Controller.

Note all the levers and other mechanisms. Current technology computers and electronic controllers have far fewer moving parts, but there is still some time required for an input change to be reflected at the output.

The bottom line is that each controller must be tuned to the system it serves to function properly, and each system will have a minimum throttling range below which system operation will no longer be stable. It is not uncommon for the proportional temperature control systems for large HVAC equipment to require throttling ranges of 5-10°F or more for stable performance. This results in large offsets or errors from the desired set point for many of the system's operating conditions. For some systems, like the one in the case study, this will result in a significant energy consumption penalty for much of the operating cycle unless the set point of the controller is constantly readjusted by the operators to compensate for the offset. Of course, most facilities groups have more than enough to do with out having to constantly run around and adjust their controllers. After all, it's supposed to be an *automatic* temperature control system. And, even if this were possible, the constant adjusting would probably result in other operational issues developing.

PID TO THE RESCUE

The good news is that the PID control algorithm (which stands for Proportional plus Integral plus Derivative) can, when properly applied, eliminate all of the proportional offset associated with a traditional proportional only control loop.

The bad news is that PID controllers are much more complex than traditional proportional controllers to tune and maintain. This problem has additional complications associated with it including:

- The exact workings and mathematics behind PID algorithms vary from manufacturer to manufacturer[ii]. The tuning constants that work for a system equipped with a controller from manufacturer A may not work on the same system if the controller is replaced by another from manufacturer B.
- The PID algorithm and how to apply and adjust it is not well understood by many in the HVAC community.
- The PID algorithm may solve significant control problems in some cases, but universal application of the control algorithm to all control loops in a facility or complex without considering if the process warrants the added complexity may cause more problems that it solves.

PID technology has actually been around in the process control industry for quite a while. The first successful PID controllers were pneumatic instruments that were developed in the late 30's and early 40's. When they were first developed, the tuning and application of the devices was actually more of an art than a science. This condition persisted until around 1940 when two engineers at Taylor Instruments – John Ziegler and Nathaniel Nichols – developed the *"Ziegler-Nichols"* method of tuning controllers[iii]. Figure 3 is a picture of a control room from a 1940's synthetic rubber plant and will give you an idea of the state of the art at the time they performed their work.

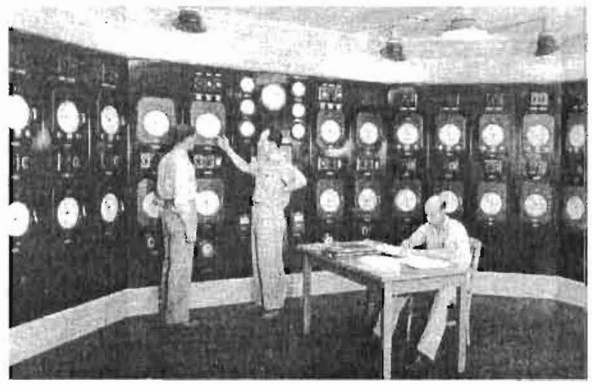

Figure 3 – 1940 Vintage Control Room

Each box with a circle in it is a PID controller. All of these controller loops and much more can now fit in a current technology PC[i].

It is the integral function (the I in PID) that really addresses the problem we are concerned with in the HVAC business; i.e. eliminating proportional offset. The derivative function allows the controller to be more responsive to changes in the process but is seldom required to any great extent in most (but not all) HVAC applications.

THE EFFECTS OF THE INTEGRAL FUNCTION

The integral function works to eliminate the proportional offset over time. There are numerous ways to actually accomplish this, and the method will vary from manufacturer to manufacturer. But, the intent in all cases is very similar. The integral function looks at accumulated offset over time (thus the term integral) and adjusts the output of the controller as required to eliminate offset. The action usually is applied in conjunction with the proportional action although integral only control is used in some limited process and HVAC applications.

Mathematically, a proportional plus integral controller will operate on some variation of the following equation:

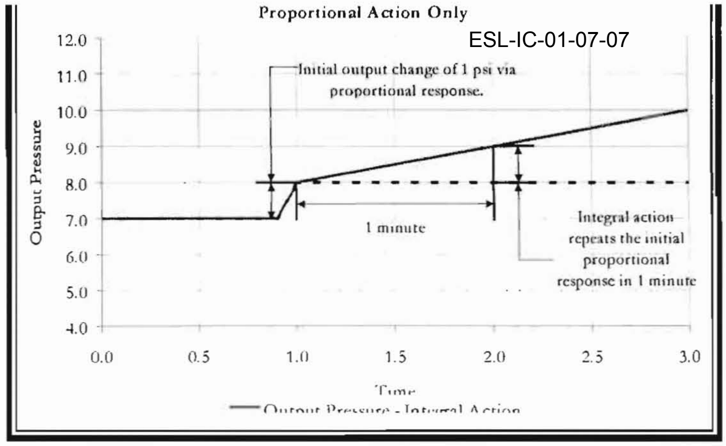

$$Output = K_P * Proportional\ Offset + K_I * \sum_{time} Proportional\ Offset$$

Where:

K_I = controller integral gain

$\sum_{time} Proportional\ Offset$ = summation or integration of proportional offset over time

As stated previously, the exact equation or algorithm used will vary from manufacturer to manufacturer, but the concept is as illustrated in the equation above. Integral gain is typically thought of in terms of repeats per minute or minutes per repeat (reciprocal terms), a reference to how quickly the effect of integral action will increase the output of the controller relative to the output change caused by the proportional response to the initial upset. This is illustrated in Figure 4.

One of the problems that seems to occur in some retrofit installations where a P only controller is replaced with a PI controller is that the once stable control loops will sometimes become unstable if the integral function is simply added to the loop without making an adjustment to decrease the proportional gain. This is because a well tuned control loop has its gain set so that the system operates just on the stable side of instability. Assuming no other changes to the control system or mechanical system it serves, anything that increases the gain of the control system significantly will force it into an unstable operating region. Thus, if integral action and gain is to be added to a well tuned proportional only control loop a concurrent reduction in the proportional gain will be required to keep the over-all system in a stable operating state. Some general rules regarding making this adjustment will be discussed in a subsequent section of this paper.

Another problem associated with integral action is called integral wind-up. Since the integral function is accumulating proportional offset over time, the accumulated value will tend to increase as long as the offset is positive (above set point) and decrease when the offset is negative (below set point). If the equipment controlled by the control loop does not have the "muscle" to eventually force the offset to be negative, the controller will simply keep accumulating positive offset, resulting in a huge accumulated value. This will drive the output of the control loop to its maximum value and it will remain there until conditions change to the point where the system can recover and force the offset to be negative *and* the negative offset has been big enough and lasted long enough to bring the accumulated positive value back down.

To illustrate this, consider a PI control loop serving a chilled water coil that is subjected to a load that it cannot handle. If everything else is at design and sized properly, the chilled water valve serving the coil will be fully open at the exact point where the coil capacity matches the load. As the load begins to exceed what the coil is capable of, the controller will begin to see a positive offset because the wide open valve can not maintain set point and the discharge temperature starts to rise. Nevertheless, the integral controller will continue to accumulate this offset with the accumulated value becoming larger and larger as time progresses. This will cause the output of the controller to reach its maximum value in short order and remain there. This will continue even when the load begins to drop back toward a level that the coil can handle, because as long as the coil is overloaded, the discharge temperature will be above set point, and thus the offset will be positive with the magnitude of the offset varying with the amount that the coil is overloaded. It is only when the load on the coil drops below what it is capable of handling at full flow that the accumulated value will start to decrease. But, since the large accumulated offset value will hold the chilled water valve wide open, it is quite likely that the system will significantly overshoot in the negative direction before the accumulated offset value is reduced to the point that the valve begins to modulate closed. The result can be sluggish and erratic system performance and wasted energy. Most controllers incorporate some sort of anti-wind-up feature, which work with varying degrees of success[iv]. Scheduled HVAC equipment is prone to this sort of problem when the system is commanded off if steps are not taken to prevent it. If the control loops remain active, they attempt to force the system to achieve set point, even though the inactive system is incapable of achieving this. When the system restarts, the loops are "wound up". One step that can help to combat this problem on programmable systems is to force the output of the loop statement to zero and/or skip the calculation step any time the component it serves is off.

THE EFFECTS OF THE DERIVATIVE FUNCTION

The derivative function responds to the rate of change of proportional offset over time. As with integral action, there are numerous ways to actually accomplish this. In the general case, the derivative function looks at the rate at which the proportional offset changes over time (thus the term derivative) and adjusts the output of the controller as required to minimize the rate of change. When properly applied, the derivative function will help to minimize the deviation from set point that a system will experience when it sees a sudden change in the requirements of the process. The need for this function is not common in HVAC systems, and thus it is often not necessary to implement it. However, the function can be useful to help minimize the swings that a system will see at start-up or due to some other large load change. Experience has show it to be particularly helpful in minimizing pressure deviations in large variable air volume systems at start-up and in dealing with the problems associated with marginally oversized valves and other final control elements.

Mathematically, a proportional plus integral plus derivative controller will operate on some variation of the following equation:

$$Output = K_P * Proportional\ Offset + K_I * \sum_{time} Proportional\ Offset + K_D * \partial Proportional\ Offset / \partial Time$$

Figure 5 - Proportional Plus Integral Plus Derivative Action with a Derivative Time of One Minute vs. Proportional Plus Integral Action Only

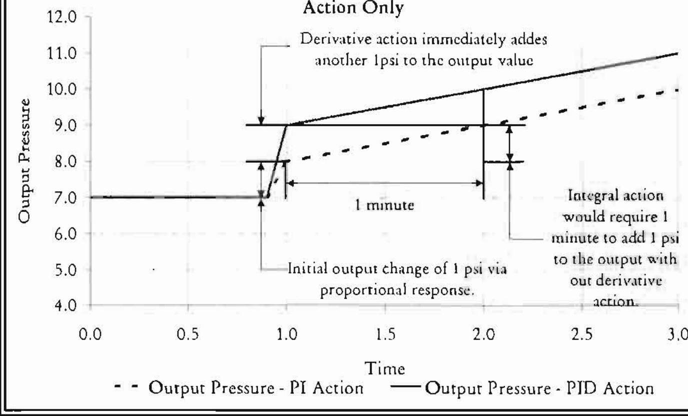

Where:

K_D = controller derivative gain

$\partial Proportional\ Offset / \partial Time$ = rate of change of proportional offset relative to time

As stated previously, the exact equation or algorithm used will vary from manufacturer to manufacturer, but the concept is as illustrated in the equation above. Derivative gain is typically thought of in terms of minutes, a reference to how long it would take the effect of integral action to increase the output of the controller in response to an upset relative to the output change caused by the immediate effect of the derivative action as a result of the initial upset. This is illustrated in Figure 5.

One of the interesting aspects of derivative action is that it only will occur during an upset when there is a change in offset relative to time. If the rate of change

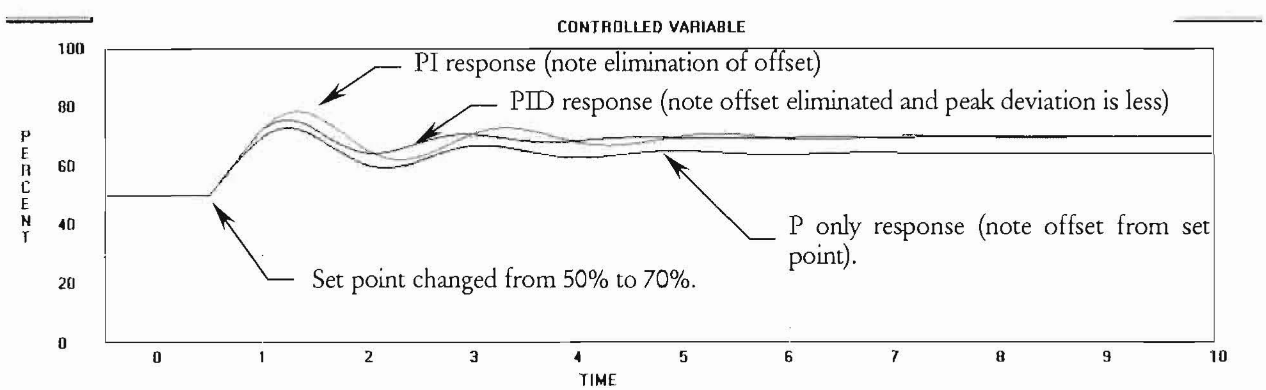

Figure 6 - The Effect of Proportional, Integral, and Derivative Control Response on System Performance

In most cases, the biggest benefit associated with PID control for an HVAC application is the elimination of the proportional offset via the integral function. Derivative action can minimize the process swing associated with a system upset, but the benefits associated with this are often quite modest or insignificant in an HVAC application.

of offset is zero, then the derivative gain is multiplied by zero and will have no effect on the output.

Derivative action, when properly applied, can result in improved performance. But, it is difficult to apply properly. With proportional, integral, and derivative action, you generally are trying "to use enough, but not too much of each function. But, with proportional and integral action, if you don't use enough, the result will generally be better than if you didn't use them at all in most cases; i.e. some benefit will be realized, but it may be less than optimal. With derivative gain, not using enough provides no real benefit, and using too much can cause many, many more problems than it cures"[v]

COMBINED EFFECT OF P + I + D

Figure 6 illustrates the response of the same system to a proportional only control loop, a proportional plus integral control loop, and a proportional plus integral plus derivative control loop. Notice how the addition of integral action eliminates the proportional offset and the addition of derivative action reduces the peak offset experienced when the loop is upset.

Eliminating the proportional offset can have significant energy and cost savings implications for HAVC systems in addition to improving the precision of the control system. In the example at the beginning of this paper, you will recall that the offset resulted in 47 tons of unnecessary cooling and 5.7 mBh of unnecessary reheat under one operating condition. Eliminating unnecessary loads like these can result in reductions in operating costs of hundreds or even thousands of dollars per year, especially in the case of large systems and systems that operate a significant number of hours per year. Using either of these functions will significantly increase the time required to tune and maintain the control loop as compared to a proportional only loop.

Minimizing the peak swing in proportional offset that will occur when the loop is upset (the effect of properly applied derivative action) provides benefits that are related more to improved operation and performance rather than improved precision and lower energy costs. Since most HVAC process are relatively steady state operations once they are stabilized and the changes that do occur usually occur gradually over a relatively long time interval, derivative action provides little additional value for most HVAC control loops.

A well tuned PID loop will exhibit a characteristic response where-in the oscillations introduced by an upset will decay fairly quickly with each peak being significantly less than the preceding. Many tuning solutions attempt to achieve a pattern called a "quarter decay ratio" in which each peak is 25% of the magnitude of the preceding one although other solutions are also considered acceptable[vi]. Generally, the decay in amplitude following a disturbance should occur in a reasonable period of time and result stable operation at set point, as can be seen from Figure 6. Loops that are stable but take a long time to settle in to set point after an upset probably require additional proportional and/or integral gain. Loops that remain in oscillation following an upset probably have too much gain.

CONTROL LOOP TUNING TECHNIQUES

A detailed discussion of loop tuning is beyond the scope of this paper. In addition, several of the sources cited contain very well written instructions and will provide excellent guidance. *Controller Tuning and Control Loop Performance, Second Edition* is especially well written and useful in that it presents the information in a non-mathematical format geared towards giving technicians and operators a practical, easily implemented understanding of PID loops and their tuning. The guide also includes the mathematics associated with the process for those who are interested. A supplemental software program is available with the guide that provides and excellent tutorial on loop tuning and allows the user to experiment with techniques and become familiar with the responses of the various types of gain. The software can also be used to quasi-simulate problem loops and play what-if games with them if you know enough about the system. Many of the illustrations used in this paper are screen captures from the output of this program (Figure 6 is one example).

In general, there are two approaches to loop tuning, closed loop and open loop. The closed loop approach is used with the system on line and operating in automatic. This is probably the approach that will be used most of the time in HVAC tuning applications. Regardless of the approach used, it is important to have some way to monitor what is going on in real time and document the results. This is especially true for the open loop method. Generally, the dynamic trending capability of current technology DDC systems can be used for this. Older systems using higher end controllers often provided this sort of monitoring in the form of a circular chart, and some stand-alone microprocessor based controllers mimic this feature electronically via a liquid crystal or CRT display. If neither of theses options are available to you, then you may want to bring a portable data-logger along that has display capabilities to use while you are testing. At a minimum, it would need to have an input that could monitor the process variable for the control loop you are trying to tune. A second input that can document the signal to the final control element is very handy, but not necessary.

The general steps for the closed loop method are as follows:

1. Turn off the integral and derivative gain.
2. Increase the proportional gain of the loop in small increments until the loop just begins to cycle. In doing this, you are finding the ultimate gain of the system, which is the point where the system is just starting to go unstable.
3. Observe the period or frequency of the cycling. This is typically called the natural period of the system and is a parameter that provides the basis for the subsequent steps and can also provide quite a bit of insight into the system's characteristics.
4. Set the controller settings to the following values.

   Proportional gain = ¼ to ½ of the ultimate gain.
   Integral time in minutes per repeat = 1.2 times the natural period.
   Derivative time in minutes = $^{1}/_{8}$ of the natural period.

5. Monitor the performance of the loop and make minor adjustments as required to optimize the performance and tailor it to the needs of the system.
6. Subject the loop to upsets by making acceptable set point changes and/or shutting down and restarting the system to be sure that stable operation is achieved in a reasonable time and with-out excessive and/or dangerous deviations in the process variable.

The open loop technique involves placing the system in manual, and when the process has stabilized, introducing a step change and observing the results. Figure 7 illustrates a typical response curve from an open loop tuning process. The general steps are as follows:

1. Place the controller in manual and stabilize the process. Its important that any changes introduced into the process during the test be changes that you made.
2. Introduce a step change.
3. When you see the results of the change, make a step change in the opposite direction with a magnitude of twice that used for the original step change.
4. Return the output to its starting value. Basically, you are trying to get the result you want and put the process back in a safe and stable operating mode with out exceeding critical system parameters or tripping safety equipment.
5. Measure and document the apparent dead time.
6. Measure and document the slope of the line in terms of rate of change per minute expressed as a percentage of transmitter span divided by the step change you introduced expressed as a percentage of controller output span.
7. Set the controller settings to the following values.

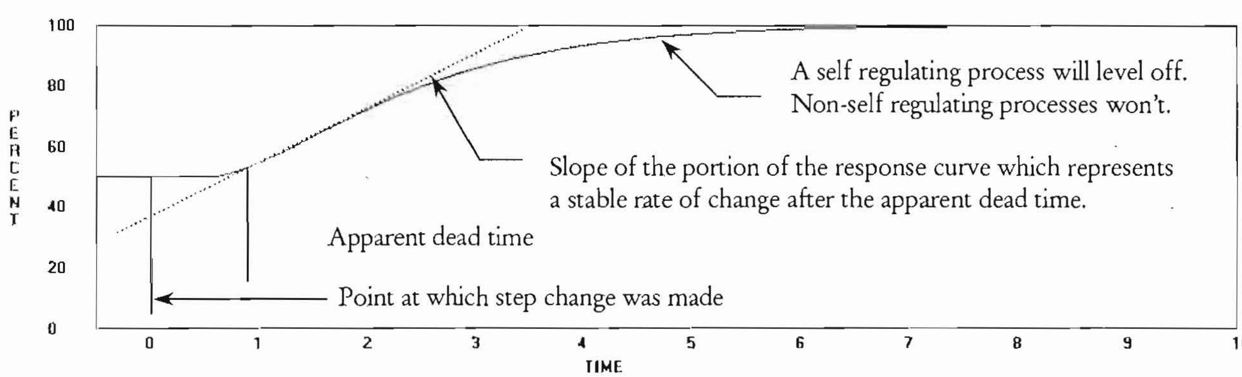

Figure 7 - Open Loop Response

The open loop tuning technique will yield a response curve similar to the following. The parameters derived from the response curve are used to determine the initial controller settings. They can also provide some insight into the characteristics of the system. Notice how the apparent dead time is made up of a flat segment of pure dead time; i.e. nothing literally happened, and a transition to a stable rate of change. In open loop tuning, you are most interested in the apparent dead time and the slope of the line where the process is undergoing a stable rate of change.

Proportional gain = 1/slope to 1/(2 times the slope)
Integral time in minutes per repeat = 5 times the apparent dead time
Derivative time in minutes = ½ the apparent dead time.

8. Return the loop to automatic and monitor the performance of the loop and make minor adjustments as required to optimize the performance and tailor it to the needs of the system.
9. Subject the loop to upsets by making large set point changes and/or shutting down and restarting the system to be sure that stable operation is achieved in a reasonable time and with-out excessive and/or dangerous deviations in the process variable.

There are some situations where these tuning process break down and won't work, but they are rare. For instance, some loops will exhibit a slope that continues to increase rather than stabilizing in an open loop test. Other loops may exhibit an reversed response to an upset. These characteristics are unlikely in an HVAC system, but are discussed in *Controller Tuning and Control Loop Performance, Second Edition* if you are interested in pursing additional information in this regard.

When tuning loops, the following general concepts should be kept in mind.

- The natural period and the apparent dead time of the control loop are very important parameters. As a general rule, the natural period will be about 4 times the apparent dead time. The apparent dead time is the result of the various lags in the system due to transportation times, slope in mechanism, thermal characteristics, etc. Generally, anything that you can do to minimize these lags will improve the performance of the control loop.
- If the control loop is subject to a periodic disturbance then the frequency of the disturbance might have a very significant impact on the loops ability to control. The control loop will be helpless in dealing with disturbances that are short relative to its natural frequency because they are too fast for it to deal with. If the disturbance is at nor near the natural frequency of the loop, then there is strong likelihood that the control loop will make things worse rather than better.
  At first you may think that repetitive cyclic upsets are not a likely situation in the HVAC industry. But, given the configuration of the systems, there usually are many control loops that interact through the dynamics of the system. If one of the loops starts to hunt for some reason, the result of it's hunting will become a periodic disturbance to other loops. Consider a VAV fan system as an example. If the system controlling the mixed air dampers starts to hunt, the dampers will start moving around, This will vary the static pressure requirements for the fan system, (especially if the dampers have not been well sized) and thereby introduce a periodic disturbance into the static pressure control loop. If that disturbance happens to fall near the natural frequency of the static pressure control loop, it could cause that loop to suddenly begin to hunt. This would lead to flow variations in the system that could impact the performance of the terminal unit flow controllers, the building static pressure control system and the control loop controlling the return fan.
- In general, for optimum performance, you want to be just on the stable side of the ultimate gain point.
- The ultimate gain of the system will change as the characteristics of the system change. In HVAC applications, this can happen for a variety of reasons including wear in the machinery and equipment and seasonal variations in the loads served and the heat transfer characteristics of the equipment used to serve the load under various conditions. Loops that are tuned in the winter may not be stable in the summer or during the swing seasons. Variable volume system loops that were tuned with the system operating at part load early in the day may not be stable at full load later in the day. You should anticipate this and expect to retune frequently during the course of the first year as the systems go through all of their operating modes in real time for the first time. In addition, you should anticipate the need to retune occasionally as equipment ages and/or systems are modified.
- Given the variable performance criteria seen by HVAC systems, you may want to be a little conservative in your tuning parameter settings, especially during the first year of operation. Otherwise, you may be faced with some major headaches when an overnight weather change takes a system that was tuned to the edge of stability over the edge.
- Before tuning, you should have an idea of what you expect to happen and what your desired outcome is as well as what the given system can be expected to do. If you have reason to believe that the valves and dampers in the system have not been sized properly, then don't expect high precision results. You should also try to get a feel for how fast the process might respond to an upset.

- You, and everyone else involved with the tuning process should know and agree about how far you will let things go before you abort the test and shut down or otherwise take over control of the system. This can be critical in some cases if you want to avoid damage to the system and equipment or avoid unnecessary safety system trips. You will need to know how to respond quickly and with confidence to a situation that is going down hill rapidly. In situations where the risks are high, it would be good for the team to rehearse their actions in the event of a run-away process to be sure everyone knows what they are supposed to do and can move quickly and with confidence.
- You should have tested all safety systems and interlocks and they should be set at the appropriate levels required to protect systems, people and equipment.
- You should schedule your testing at a time when the system and the loads it serves can tolerate some disturbance and a shutdown down of the system would not create a crisis. Tuning the discharge temperature control loop on a surgery air handling system during an open heart case is not a good idea.
- Plan on being readily available for some time after you make your adjustments in case problems show up. Making significant changes to a control system late Friday afternoon on your way out of town on a two week vacation will not make you very popular upon your return if the changes result in operational problems and you aren't there to help correct them or return things to the way they were.
- Document everything. This includes the settings that were in placed prior to starting the process, the settings you left in place when you finished, and any other pertinent observations that you might have made during the process. This information should be communicated to the folks who will be running the systems right after you are done tuning. If there are latent problems related to the final tuning parameters you left in the system, they will often (but not always) show up in the first hours or days of operation subsequent to your adjustments.
- Proceed slowly. The best thing that you can do after you make an adjustment to a tuning parameter is wait and watch. Often, it is tempting to make a second change fairly quickly if the initial change you made doesn't produce the anticipated result. This may be a satisfactory approach if you are very familiar with the system and the process is not particularly critical. However, bear in mind that the lack of anticipated response could be due to system lags or other phenomenon that you failed to consider and you may suddenly find yourself suddenly dealing with a run-away process when the accumulated impacts of your changes finally take effect.

RELATED ISSUES

There are several control loop issues that are not directly related to the actual tuning process but which will certainly impact it. These include:

- *Non linearity* - Non linear characteristics abound in the process and equipment that is associated with HVAC systems. Most heat transfer devices have a non-linear relationship between heat transfer rates, flow rates and temperature differences. Many of the sensing systems that provide inputs to the controllers we use are non linear. Examples include the output from the thermistors commonly used to measure space temperature and the output from differential pressure based flow measuring elements. Velocity limiting is another non-linear system response that can cause havoc with HVAC. Actuators and final control elements generally have some finite speed at which they can move through their full stroke. If the process goes through a large change, and the rate of change is faster than the rate at which the final control element can move to make a correction, then the process becomes velocity limited. If the change in the process is small, the problem usually doesn't show up. As a result, the loop may be stable for small upsets but unstable for large ones.
- *Loop interactions* - Sometimes, designs include interactive loops where in the normal response of one loop will upset the second loop, which will upset the first loop, etc. A pipe line with a pressure control loop that modulates a valve to maintain the pressure ahead of a second valve that is controlled by a flow control loop is a good example of this. Many times, one of the control functions can be eliminated to solve this type of problem. Or, one of the loops can be tuned to be very "loose" while the more critical loop is tuned for "tight" control.
- *Auto tuning* - Auto tuning is one of those things that sounds like a really nice feature, but in fact can be of little use or benefit in some situations. It is probably unwise to rely on it as a total loop tuning solution, despite what sales literature and salesmen may lead you to believe. The exact auto tuning algorithm employed will vary from manufacturer to manufacturer. There have been instances where an auto tune controller from manufacturer A would not be able to tune itself in a certain process, but when a controller from manufacturer B was substituted, the same process was able to stabilize. In a different process, the controller

from manufacturer A was able to self tune, but the controller form manufacturer B only achieved marginal results. It is also important to understand exactly how the auto tune algorithm works. Some algorithms make gradual adjustments based on ongoing observation of the process. Others will actually do a variation of the open loop method and slam the final control elements up against their limits. This can be an undesirable approach on some processes.

- *Hysteresis* – Packing friction and play in a linkage system can cause control problems when small errors exist, especially where integral control is in effect. Basically, the controller tries to respond to a small change, but nothing happens due to hysteresis. So, the controller increases its output some more, which causes too much to happen and the cycle repeats itself in the reverse direction. The characteristic indication of this problem is a low amplitude cycle in the process. If you respond to this problem based on the normal tuning rules, it won't be solved; it will only change the period of the cycle. Thus it is important to recognize and distinguish this problem form hunting associated with marginal stability. One very practical way to check for this it to simply place your fingertip on the valve stem near the packing. You can often feel the stem "popping" back and forth when the output from the controller over-comes the packing friction (as opposed to a smooth modulation.) Or you may be able to observe or feel actuator motion with out a subsequent valve motion due to play in a linkage system. In any case, solving this problem requires making changes that minimize the hysteresis rather than additional controller tuning.
- *Matching final control element spans to output spans* – If the output span from the control system is not matched to the actual span of the final control element, a problem similar to wind-up will occur because the control loop spends a portion of its time trying to actuate something that can not be further actuated. Then, when things change, the controller spends time backing down from its limit to a point where the final control element begins to actuate again.
- *Filters* – Because current technology controllers use digital technology and operate at very high speeds, they are capable of detecting and responding to very small and inconsequential changes in the process. This can, in human terms, drive the control system crazy trying to respond to things that really don't matter. It can also ruin certain types of actuators in a matter of months. As a general rule, the filter time for the system (if available) should be set only as long as is necessary to provide the desired filtering, regardless of whether the loop is P, PI, or PID.
- *Sampling rate* – This issue is related to the filtering issue we just covered. The sampling rate used by a controller to sample and control the process as well as by the engineer to monitor the process can have a tremendous impact on what you (or the controller) perceive as actually going on. Figure 8 illustrates this phenomenon, which can mask and distort the information that is presented to the control system and operators.

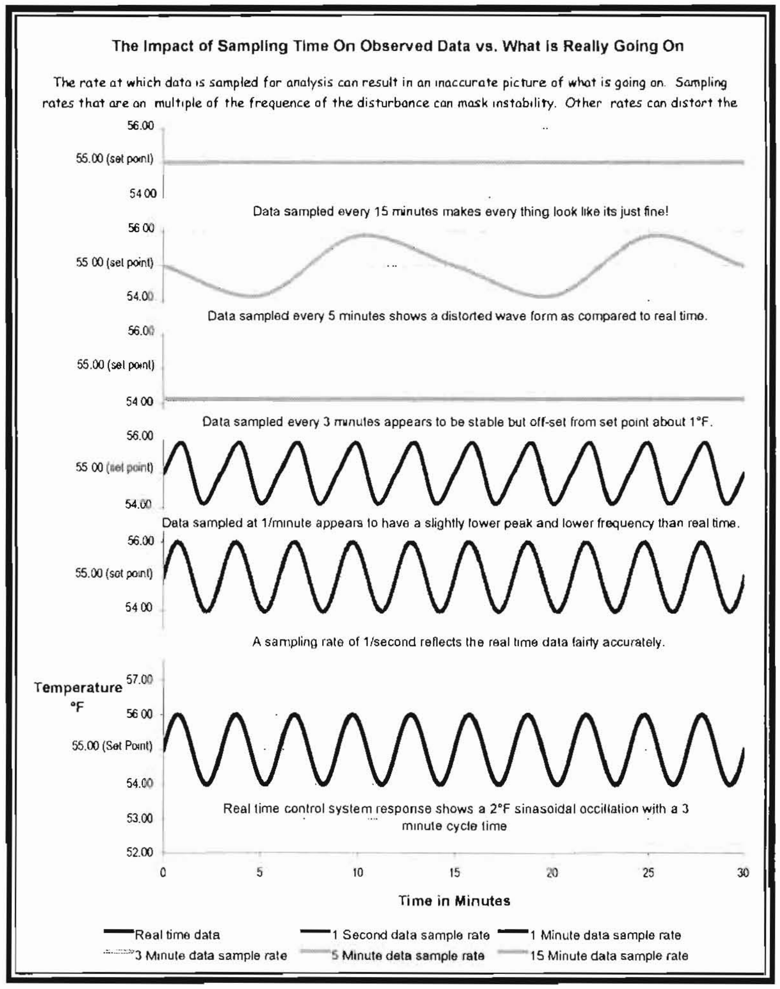

WHEN TO APPLY PID INSTEAD OF P ONLY

It is important to understand that the benefits associated with integral and derivative action come at a price. Just because you can do PID doesn't mean you should do PID. Invoking the integral and derivative functions will result in loops that require more time and attention to properly tune them. In addition, these loops will require more attention over the course of their operating life to maintain and adjust the tuning parameters as the system and load characteristics change. In addition, the operators charged with this function will need to have a more sophisticated understanding of control theory and its application in a real world operating environment. The following general guidelines are suggested when making decisions regarding the application of PID to a given control requirement.

- Use proportional only control in situations where high precision is not required or warranted by operational or economic concerns. A prime example of this is zone temperature control. In quite a few instances, a simple, well tuned and calibrated P only control loop will provide very satisfactory control for space temperature. Evidenced of this can be found by recognizing that numerous buildings to this day use pneumatic thermostats for zone temperature control with satisfactory results. In general, most of the problems with this approach can be traced to low end equipment (one pipe thermostats for instance), poor maintenance and calibration, and thermostat locations issues (over the coffee maker for instance). Integral and derivative action will do nothing to solve these more fundamental problems and will significantly increase the initial set up and maintenance time for the system.
- Another example where a proportional only loop might prove to be satisfactory is for secondary back-up or limit applications. Mixed air low limit control loops are good examples of this.
- Cascaded or highly interactive control loops are another area where application of PID to every single loop may yield more problems than satisfaction. It may be better to use PID for the critical loop and allow the other loops to function as P only loops.
- Add integral action in situations where precision is required. Controlling chilled water temperatures or building static pressures are good examples of this type of application. In these situations, minor offsets from set point can have significant operational and energy issues associated with them and eliminating the offset via integral action will provide significant benefits.
- Add integral action in situations where the proportional offset associated with a proportional only loop will result in significant energy waste. The reheat fan system example used at the beginning of this paper is a good example of this type of situation.
- Think hard before adding derivative action to a control loop. To be effective at all, it must be very carefully applied and adjusted. If implemented improperly, it can cause many more problems than it solves. Generally, HVAC systems can be made to perform quite well with out the use of this function. It can prove to be beneficial in situations where systems experience significant deviations in the process parameters at start up and in dealing with final control elements that are marginally (not significantly) oversized.

---

[i] Photo taken from *Pneumatic Instruments Gave Birth to Automatic Control, from Reference Guide to PID Tuning, A collection of Reprinted Articles of PID Tuning Techniques*, reprinted from Control Engineering Magazine, Cahners Reprint Services - 800-323-4958 ext. 2240, page 33

[ii] See *The PID Algorithm for the Process Industries* at http://members.aol.com/pidcontrol/pid_algorithm.html.

[iii] *Modern Control Started with Ziegler-Nichols Tuning, from Reference Guide to PID Tuning, A collection of Reprinted Articles of PID Tuning Techniques*, reprinted from Control Engineering Magazine, Cahners Reprint Services - 800-323-4958 ext. 2240, pages 13 through 15 contains an interesting interview with John Ziegler where-in he talks about the development of the controller and it's tuning techniques.

[iv] *Controller Tuning and Control Loop Performance, Second Edition*, page 9, David St. Clair, Straightline Control Company Inc. 3 Bridle Brook Lane, Newark, DE 19711-2003,Dwstclair@aol.com, http://members.aol.com/pidcontrol/.

[v] *Controller Tuning and Control Loop Performance, Second Edition*, pages 11-12.

[vi] *A Comparison of Controller Tuning Techiques, from Reference Guide to PID Tuning, A collection of Reprinted Articles of PID Tuning Techniques*, reprinted from Control Engineering Magazine, Cahners Reprint Services - 800-323-4958 ext. 2240, page 18.