29: PID Monitoring and Diagnosis

Myke King describes how PID controllers might be monitored and how to diagnose common faults

THERE is a growing range of products available to monitor and report on PID controller performance. To differentiate their product from the competition and justify quite costly licences, vendors include rich functionally.

But many products report parameters that have little value or are too complex to be fully understood by the user. While such a product may ultimately be justified, the user would do better by first developing in-house tools. The data are available (or can be readily made so) in the control system historian and simple reports can be built using, for example, Excel. Features can be extended over time, driven by the engineer's needs.

Typical candidate parameters are outlined in Table 1. If, at a later date, it is believed a commercial product is justified, the engineer will have a valuable benchmark against which to assess those offered.

QUICK READ

- **Simple, in-house monitoring can be highly effective:** Useful PID performance insights can be generated from historian data using basic tools, providing a strong benchmark before investing in commercial software
- **Valve and hardware issues often dominate poor control:** Stiction, hysteresis and nonlinearity in final control elements are common root causes of PID problems and can be diagnosed from characteristic trends
- **Performance logging enables better decisions:** Systematic monitoring of controllers and analysers helps distinguish tuning issues from equipment faults and supports evidence-based maintenance or replacement

Table 1: Example PID performance metrics

CATEGORY	METRIC	PURPOSE
Plant-wide indices	Fraction of controllers in automatic	Indicates overall level of control automation
	Total number of controllers in alarm	Highlights systemic control or process issues
Controller status (percentage of time)	In manual mode	Identifies controllers not operating under closed-loop control
	In alarm	Flags persistent or nuisance alarms
	Saturated (output <0% or >100%)	Indicates poor tuning or actuator limitations
	Measurement failed validity check	Identifies instrumentation reliability issues
Process and hardware limits	Min/max measurement range	Assesses whether transmitters are correctly ranged
	Min/max valve position	Indicates potential valve oversizing or undersizing
Control quality indicators	Average absolute error or standard deviation	Assesses recent tuning or performance improvements
	Maximum deviation from setpoint	Highlights worst-case control excursions
Capacity and interaction	Use of surge capacity (eg averaged level control)	Indicates buffering effectiveness
	Measurement comparison (eg inferential vs lab)	Identifies bias or analyser issues
Event-based metrics	Number of alarms	Measures alarm frequency
	Time spent in alarm	Highlights severity and persistence
	Mode changes (auto/manual)	Indicates operator intervention frequency
	Retuning events	Flags unstable or poorly performing loops

Figure 1: Audit of controller problems

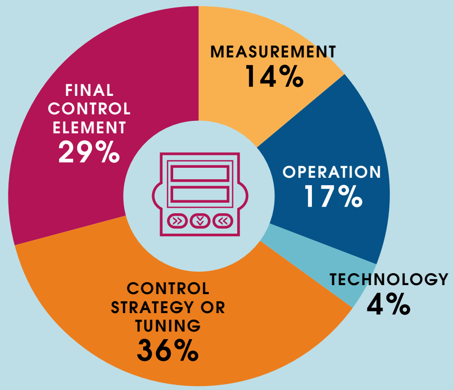

Figure 2: Effect of stiction on open loop behaviour

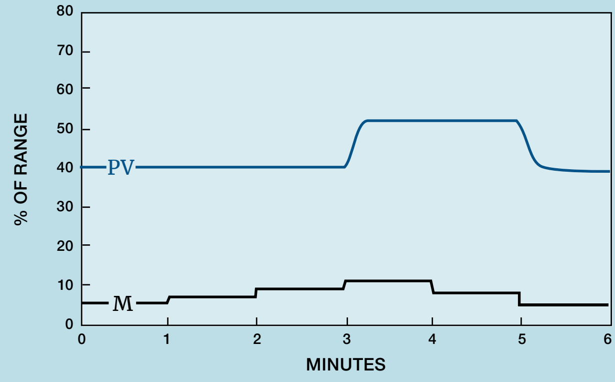

CONTROLLER FAULTS

Figure 1 shows the result of an audit conducted by Emerson, one of the leading control system providers. While poor control strategy design and its tuning is the largest cause of performance problems, we have dealt with these in previous articles. Here we focus on the next biggest cause – issues with the final control element, usually the control valve. While maintenance of instrumentation normally falls outside the remit of a chemical engineer, the ability to identify the source of a control problem is important.

CONTROL VALVE STICTION

A combination of the words "stick" and "friction", *stiction* is the force necessary to start movement of the control valve stem. It is the threshold required to overcome static friction. In a poorly lubricated valve, this can be well above the force required to overcome dynamic friction. The problem can develop long after the controller was commissioned.

Figure 2 shows the behaviour of a controller during open-loop step-testing. The small changes in controller output (M) have no effect until a sufficient change is accumulated. The valve then responds to the accumulated change to give the increase in the PV. Normally such testing would permit identification of the process gain:

$$K_p = \frac{\Delta PV}{\Delta M}$$

With stiction, it is unlikely that a reliable result will be obtained.

Figure 3 shows the behaviour in closed-loop mode. While there are several causes of such oscillation, most indicate that the controller gain (K_c) should be reduced by a factor of around 2. Figure 4 shows the impact of doing so. The unexpected increase in amplitude, and reduction in frequency, is typical of stiction – and so gives us a means of diagnosing it.

Figure 3: Effect of stiction on closed loop behaviour

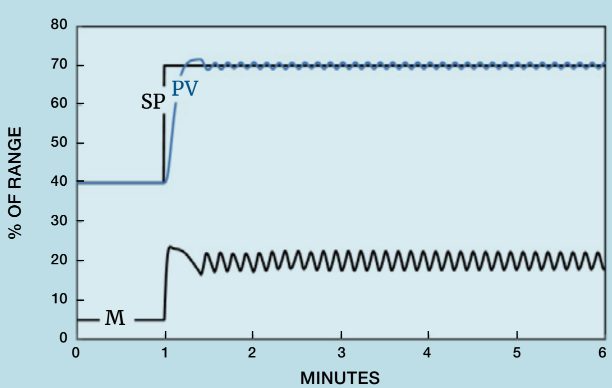

Figure 4: Diagnosing stiction by reducing controller gain

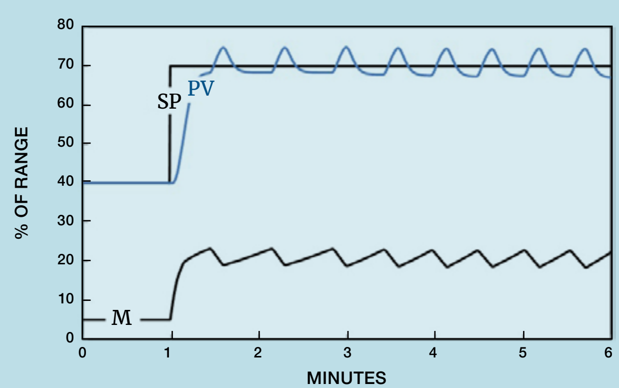

Figure 5: Explanation of hysteresis

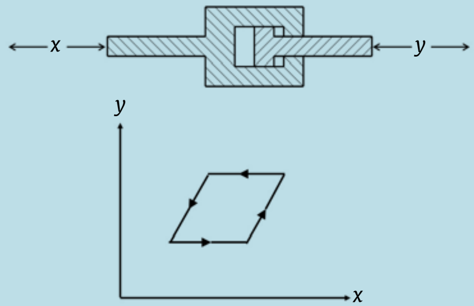

CONTROL VALVE HYSTERESIS

Hysteresis is the dependence of valve movement on previous changes. Also known as *backlash*, it is caused by slack in the valve positioner's moving parts. This can develop as the parts wear over time. Figure 5 illustrates the behaviour. Representing movement of the valve, changes in the signal to the valve (x) are not exactly replicated by movement of the valve stem (y). Whether an increase in y will follow an increase in x will depend on whether the last move was an increase. Figure 6 illustrates the impact hysteresis has on step-testing. Each of the steps made to the controller output are the same size but the resulting change in valve position varies depending on the direction of previous moves. As a result, the process gain could not be reliably estimated.

Figure 7 shows the impact that hysteresis has on closed-loop control performance. While the approach to setpoint appears reasonable, the controller output shows more erratic behaviour. Figure 8 shows a real example, where the controller output is unnecessarily changing by ±20%. In this case it was possible to also plot flow against controller output, shown as Figure 9. The cloud of points cover several days, with one particular track highlighted. This shows the parallelogram trend typical

Figure 6: Effect of hysteresis on open loop behaviour

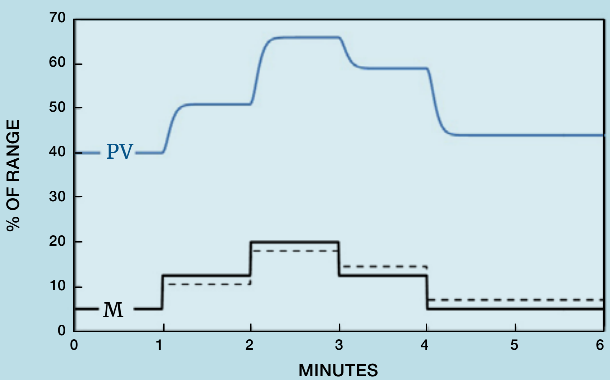

Figure 7: Effect of hysteresis on closed loop behaviour

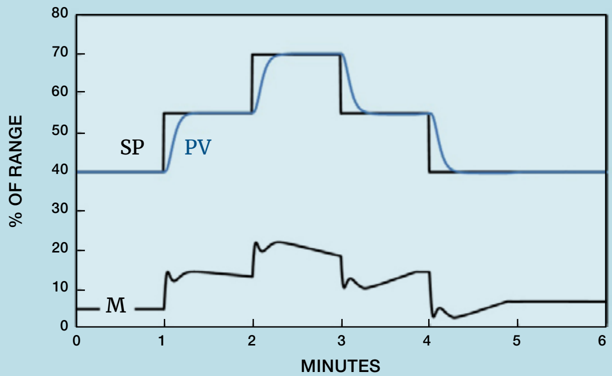

Figure 8: Real example of hysteresis

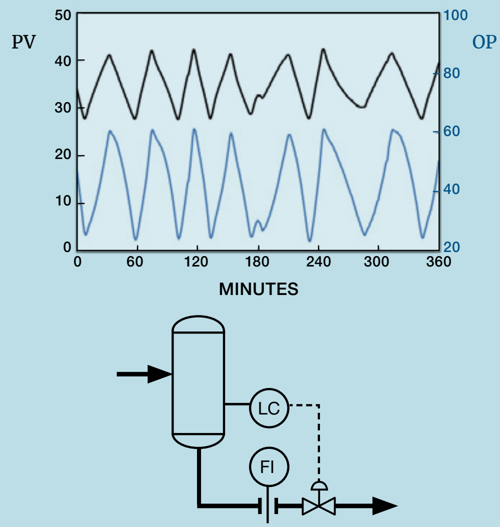

Figure 9: Diagnosing hysteresis

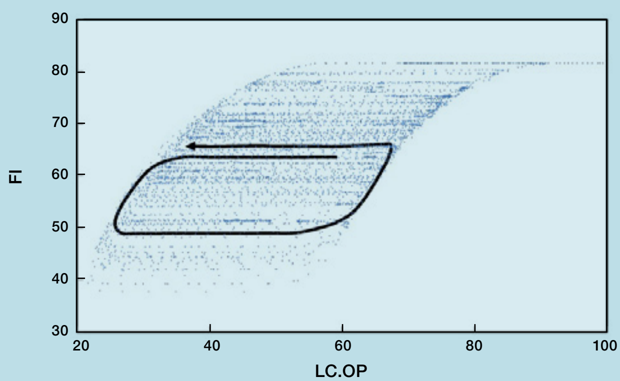

Figure 10: Effect of nonlinearity on closed loop behaviour

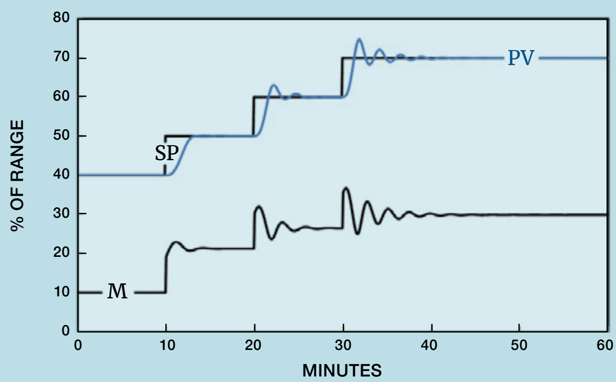

of hysteresis. Showing particularly poor performance, unnecessary ±15% changes are made to the flow. Overhauling the control valve resolved the problem.

NONLINEARITY

Figure 10 illustrates the problem. While the controller responded well to a setpoint change from 40 to 50, it becomes increasingly oscillatory as the setpoint is increased further. The process gain is clearly greater at the higher values. In *TCE 990/991* we covered techniques aimed at removing, or at least reducing, nonlinearity in a level controller. In *TCE 995*, we similarly covered flow control, where the PV (the flow) is zero when the MV (the valve position) is zero. We also showed how a change to an equal-percentage valve type might be applicable to, for example, a temperature controller manipulating a bypass around a heat exchanger. But in the situation where the wrong valve has been installed, signal conditioning provides a more immediate and potentially permanent solution. Further, there will be cascade controllers where the measurement of the primary controller is not linearly related to the setpoint of the secondary.

Figure 11 takes the steady state conditions from Figure 10. It shows that the process gain, between the extreme operating regions, varies by ±53% – well above the ±20% we

Figure 11: Variation in process gain

PV	*MV*
40	10.0
50	21.2
60	26.4
70	29.9

$$K_p = \frac{10}{11.2} = 0.89$$

$$K_p = \frac{10}{5.2} = 1.92$$

$$K_p = \frac{10}{3.5} = 2.86$$

Figure 12: Linearised PV

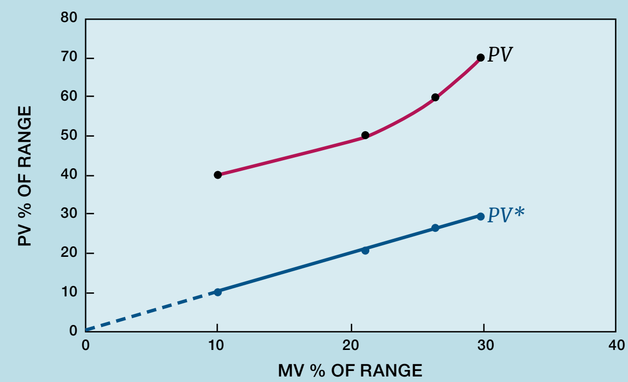

typically tolerate. To resolve the problem, we choose a nonlinear function, for example:

$$PV^* = a_0 + a_1PV + a_2PV^2$$

We require:

$$PV^* = K_pMV + bias$$

But we can choose:

$$K_p = 1 \quad \text{and} \quad bias = 0$$

$$\therefore MV = a_0 + a_1PV + a_2PV^2$$

Fitting this function to the three data points gives:

$$PV^* = -69.6 + 2.77PV - 0.0193PV^2$$

Figure 12 shows the effect of this signal conditioning on the relationship between the PV and the MV. To make it transparent to the process operator, the same calculation is applied to the setpoint. Figure 13 shows the impact of the modification.

Figure 13: Nonlinearity resolved

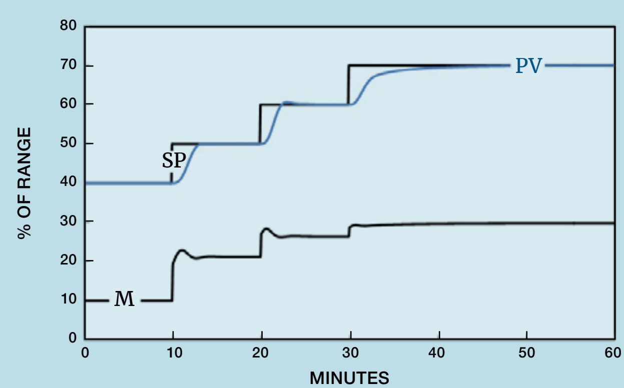

Although not perfect, there is greater consistency in the responses. If required, fitting a higher order function would give further linearisation.

NOISE

In *TCE 994* we described a range of noise reduction techniques. A common problem is filtering being unnecessarily included in a controller. Usually the criterion is what "looks good" for a trend of the PV. Depending on the controller tuning, the noise transmitted to the control valve can be amplified or attenuated. It is the amplitude of this that should determine the need for filtering. Monitoring this aids in tuning the filter and, perhaps more importantly, exploring the effect of increasing the scan interval. There are also simple tools available, based on Fourier analysis, that enable a filter to be designed that removes a specific range of frequencies from the noise spectrum.

OSCILLATION

Considerable research has been completed to develop techniques that detect a controller is oscillatory. Some such techniques also aim to identify which controller is the underlying cause. One might argue that a simple visual check of the PV trend might achieve the same, with potentially offending controllers sequentially switched to manual to identify the cause. However, there are examples, particularly with average level control, where very slow oscillations (caused by poor tuning) are mistakenly assumed to be the controller dealing with flow disturbances.

ON-STREAM ANALYSERS

While analyser reliability has improved over the years – and there is now greater recognition of the importance of good sample system design – analysers can still be problematic. It is therefore understandable that process operators are cautious about placing them in automatic control. An undetected failure can be extremely costly: there are many examples of a product tank only being discovered to be off-grade when a routine laboratory sample is taken after the tank is already full. Resolving such incidents is disruptive and expensive and the costs can far outweigh the quality improvements the analyser was intended to deliver. Further, even if the analyser is fully repaired, there will be a reluctance to again use it for automatic control.

Analyser monitoring should be designed for two purposes. The first, of course, is to rapidly detect a failure and then disable control. The second is to record performance over time and so present evidence that, after such a catastrophic failure, it is now reliable. Or, if it remains unreliable, be able to present a proven argument that a replacement is justified.

There are several parameters that can be monitored. The first and simplest is a HI/LO check on the measurement. The HI and LO settings need to be as close to the true value as possible without generating spurious alarms. Next, a properly set limit on rate-of-change would identify a measurement changing faster than the process dynamics would permit. Discontinuous analysers, such as chromatographs, will, after a failure, continue to transmit the last good value. If not detected then the controller will ramp the process conditions, until some operating constraint is reached, to eliminate any deviation from setpoint. Failure is best detected by monitoring the time interval between measurement updates and alarming if this goes beyond what is expected. Such analysers have a "read-now" contact that can readily be connected to the DCS. Finally, if the analyser is supported by an inferential property (*TCE 1004–1007*), an excessive deviation between the two would indicate failure of either.

In the event of detection of failure, normal practice would be to disable to controller and issue a warning to the process operator. If the analyser comes good again, then the procedure for recommissioning the controller needs to be agreed. For example, if the outage is short, this might be automatic. If longer term, then operator intervention might be required.

Analyser performance should be logged over time. This should include time-stamped events. Additional tags should be added for each. These are set to 1 if the analyser is good, 0 if not. Averaging these tags allows, for example, monthly service factors to be reported, to prove reliability or justify improvement. ■

While analyser reliability has improved over the years… analysers can still be problematic. It is therefore understandable that process operators are cautious about placing them in automatic control

NEXT ISSUE

In the next issue we will address the importance of effective training in the practical aspects of process control.

Myke King CEng FIChemE is director of Whitehouse Consulting, an independent advisor covering all aspects of process control. The topics featured in this series are covered in greater detail in his book Process Control – A Practical Approach, published by Wiley in 2016