Observing the Data

There is frequently a tendency to 'leap in' and tune control loops regardless of the actual problem. It is very often the case that process instability has its root cause in operational or instrumentation issues rather than inadequate tuning. So the very first stage in approaching a tuning exercise is to observe the data. This article introduces the importance of collecting historical trends with sufficient resolution to enable fault-finding and process-learning. Detailed analysis at this stage may indicate that the problem is not a tuning issue at all.

High Speed Trending

Controller tuning is often attempted using the trend facilities available on the operator consoles, which are typically updated at relatively slow speeds. Extreme care needs to be taken when using such trends to ensure that vital information is not being overlooked due to the low update times. For example, Figure 1 shows the level input and valve command of a standard level controller taken from a one-minute trend display. In the first minute, the valve has closed in causing the level to increase. In the second minute, the controller reacts by opening the valve causing the level to drop again. In the third minute, the level is too low so the valve closes in. It appears that the level controller is cycling on a one-minute period due to the command 'fighting' against the level, possibly due to valve stiction or too much integral action.

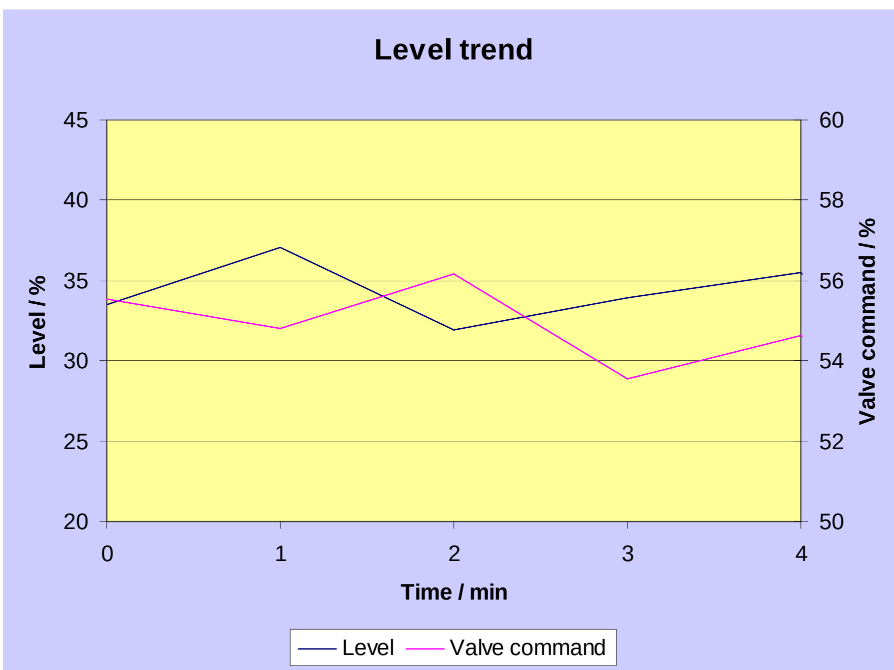

Figure 1. Low sampling rate

Figure 2 shows the same level controller but with the level trend at a much higher frequency of 4 samples per second. With the increased detail, it is now clear that the actual level measurement includes a significant noise component (about 8% of the full instrument range). The low speed trend is sampling the actual data every minute and, dependent on the noise when the reading is taken, can be anywhere within 8% of the real value. Hence, the reason for the apparent 'fighting' between the valve and level on the slow trend. In fact, comparing the actual level and the actual command trends on the high speed display, the command is more or less following fluctuations in the level, which were due to unstable inlet flows. The primary issue is measurement noise and not poor tuning.

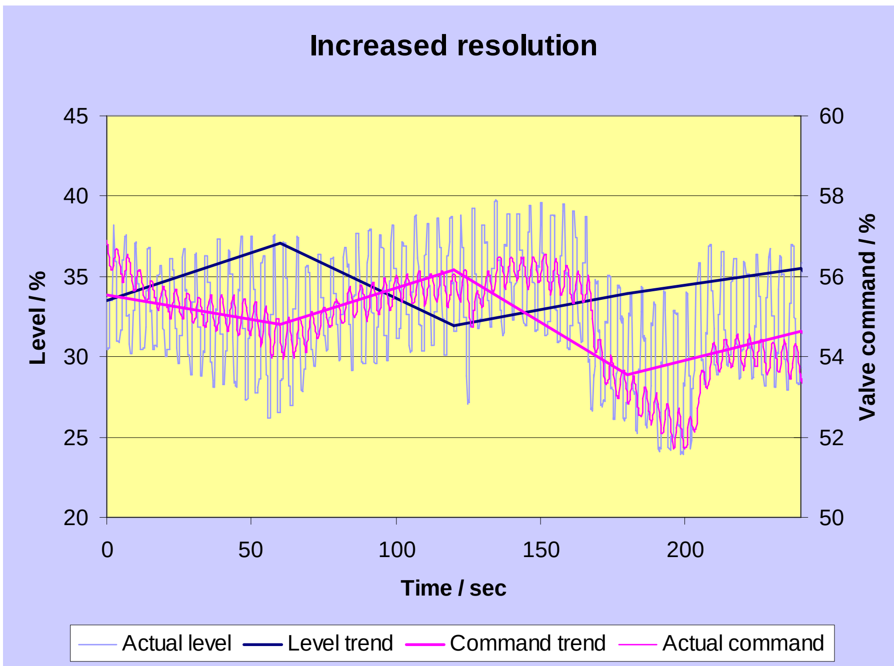

Figure 2. Low sampling rate

The main advice here is to analyze the controller performance using trends with sufficiently high update times to show the details of the measurements. With slow responding level controllers, an update time of 5 seconds is sufficient for tuning purposes, but may not show up noise- or valve-related problems so a high speed is preferable. For fast-acting loops such as compressor anti-surge controllers, the 5-second update time is far too short and needs to be in the order of 0.05-0.10 seconds to capture the sensitive process details when approaching surge conditions.

There is generally a trade-off on most control system trend displays between the trend update time and the trend time span: a 5-second update time may mean that only 5 minutes of history is displayed on the console, which is far too short for adequate tuning of slow loops. Loop analysis is much easier when the complete response is available on one display and even better if the time span can be extended to allow comparisons between two or more previous responses. This may mean a time span of one or more hours of 5-second trend for slow loops, but only five or ten minutes of high speed trend for fast loops.

By far the best solution is to invest in data acquisition equipment which can be connected to any control system input or output. The main advantage is that the signals are measured directly prior to any filtering, sampling or processing by the control system. Several manufacturers produce PC-based data acquisition packages with a box of electronics communicating with the PC through serial or USB links. Ensure that the electronics have suitable isolation between each channel and between the channels and the PC to prevent shorting out the control system. The trend data in this article has been collected using Dataq Instruments and Techmation Protuner equipment.

Separator Level

The fluids coming from an oil well are a mixture of oil, water and gas. A separator is a vessel used to separate oil, water and gas: the liquids settle out with the oil floating to the surface of the water and the gas bubbles off. A 'weir' placed across the vessel creates a barrier so that the water and oil can flow out through different outlets. There are three controllers on the vessel, one controlling the pressure, the second controlling the oil level and the third controlling the water level.

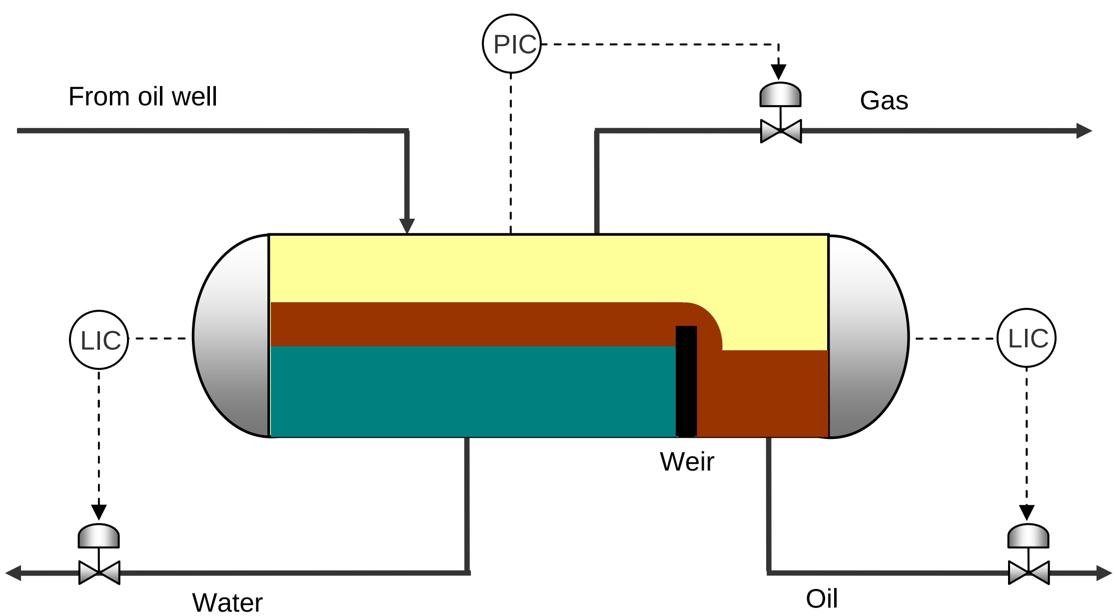

Figure 3. Separator control loops

Figure 4 shows the oil level in a separator together with the command from the oil level controller to the level control valve. The level is moving up and down by as much as 30% every couple of minutes. This particular plant had experienced major instabilities for many years, causing considerable workload for the operators and continually risking inadvertent shutdowns. Notice how the valve command is also winding around. This causes the oil flow from the separator to change continually, thereby introducing instability into the next vessel downstream.

The reason for the instability is clear by observing that the rate of the valve command changes almost every time the oil level crosses 60%. Figure 2 shows further data on a shorter time scale and this time both the level and the command change the rate of rise when the level reaches 60%. The oil compartment in this separator was rather small so that small fluctuations in the incoming or outgoing flow result in large level swings. Since there was plenty of space above the weir, the setpoint of the oil level controller was increased to operate the vessel in 'flooded' mode.

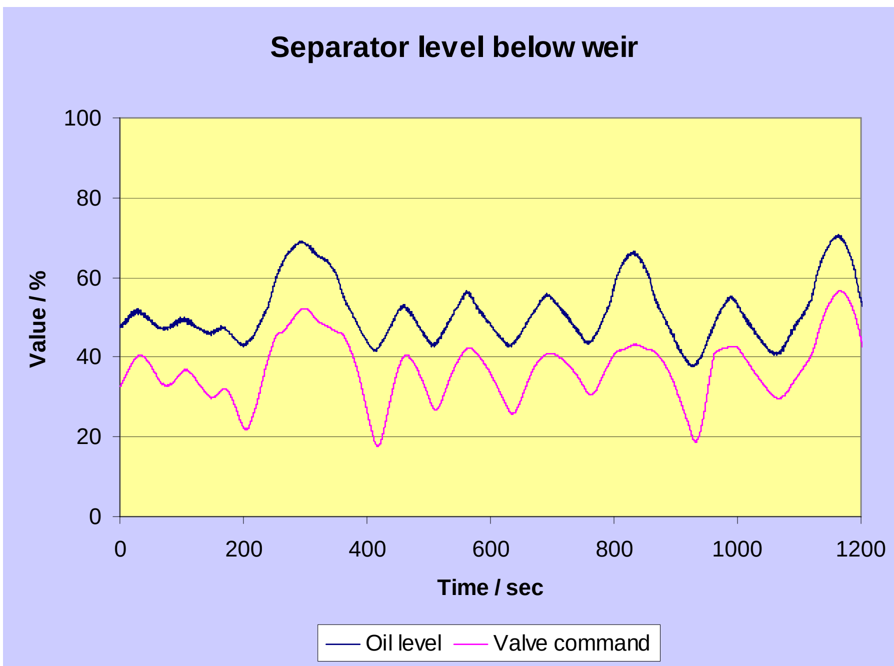

Figure 4. Separator oil level and control valve output

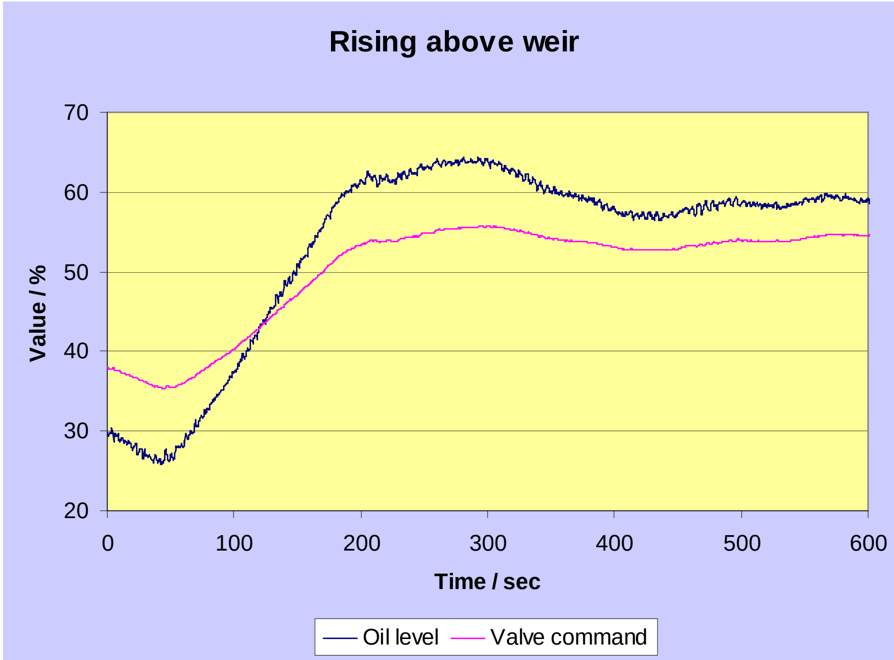

Figure 5. Separator oil level rising above weir

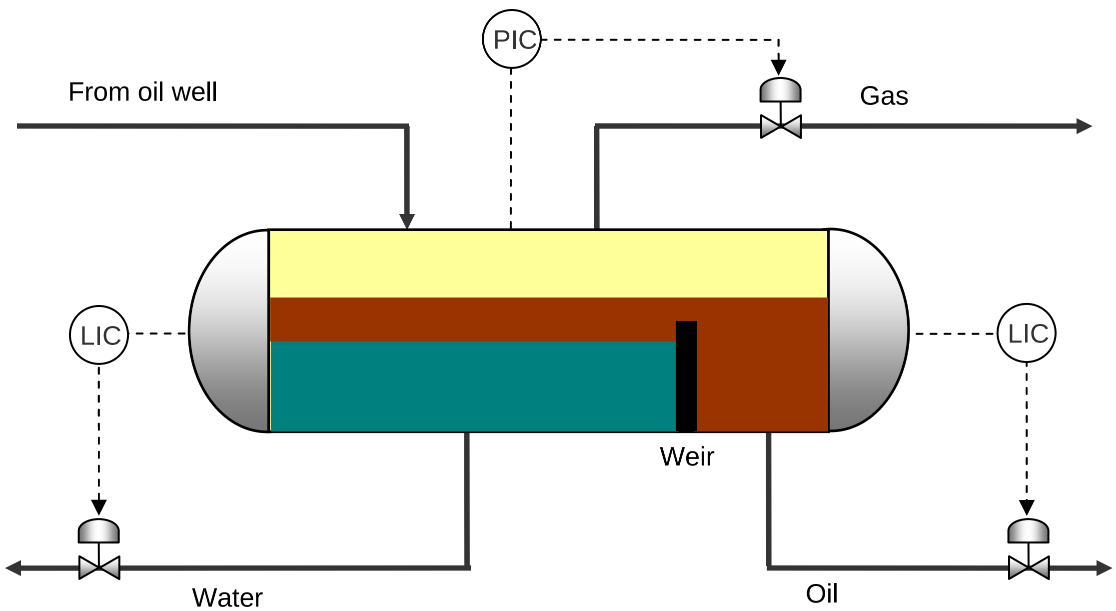

Figure 6. Flooded weir operation

The oil now has a larger volume over which to dissipate any flow fluctuations and, as shown in Figure 7, the oil level is only swinging by about 10% and over a much longer period. Since the oil control valve does not need to move around as much, the flow from the separator to the downstream process is much smoother and allowed much of the plant to be tuned up correctly.

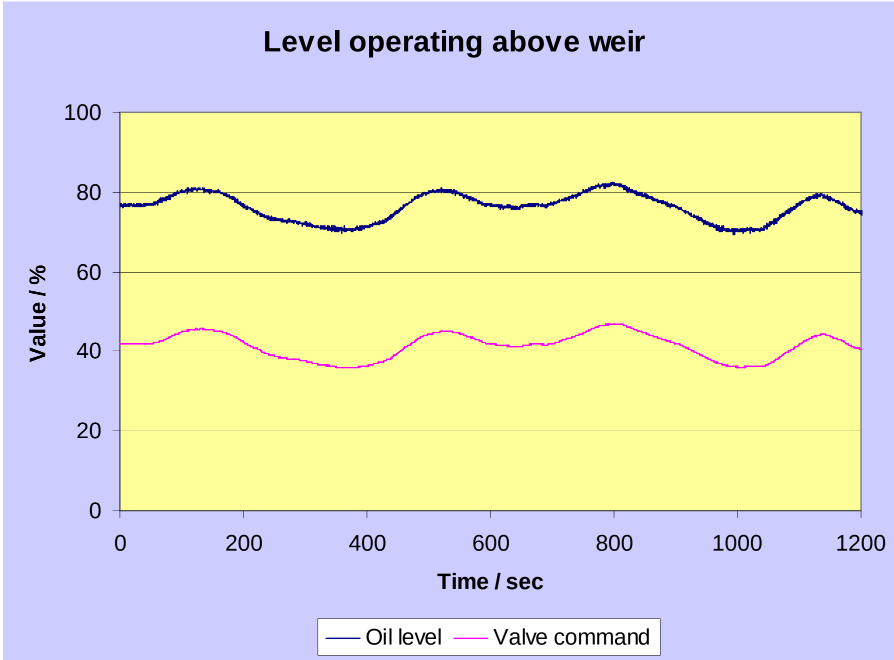

Figure 7. More stable oil level

Data Across Network

Figure 8 shows the pressure input and valve command of a differential pressure controller during a small setpoint change. The command is changing in a step-wise manner and starts to cycle after the differential pressure has reached the new setpoint (from about 220 seconds onwards).

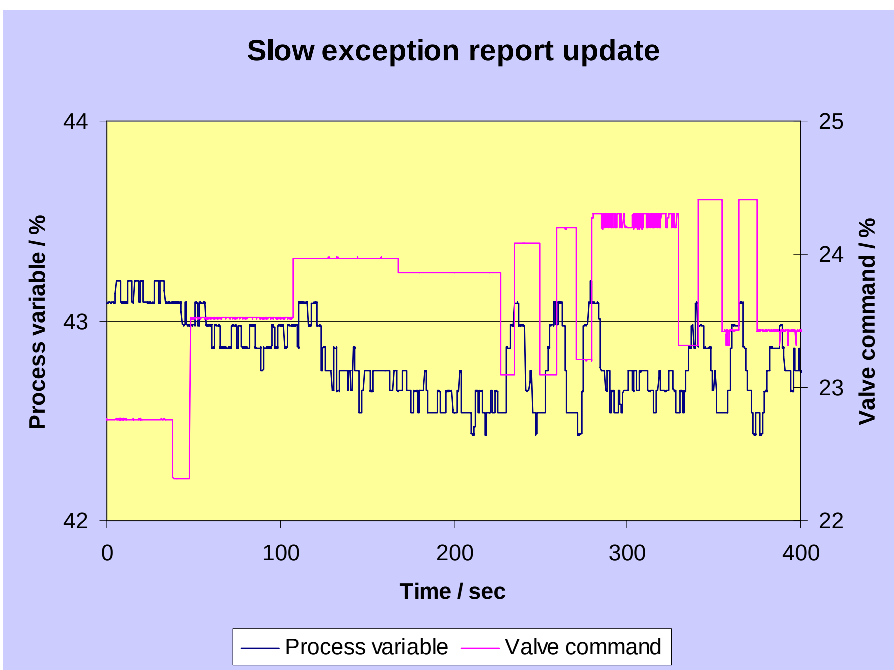

Figure 8. Differential pressure controller

This particular controller had been configured such that the valve command was transmitted over a network to another part of the control system before being copied to the output signal to the valve. In order to reduce loading on the network, the valve command was only transmitted if:

- The command changed by 1% or more.
- One minute had elapsed since the last transmission

This can be seen quite clearly in the trend display. From 50 to 200 seconds, the command does not exceed a 1% change so the new value is only transmitted every 60 seconds. After 220 seconds, the command steps closed by 1% causing the differential pressure to rise. Since the differential pressure is now too high, the controller commands the valve to open again resulting in too large a drop in pressure and so commences the cycling.

The solution was to change the 'transmission' rule to send the command on changes of 0.01%. Figure 9 shows the response to a setpoint change, now much smoother and not exhibiting the same degree of cycling as previously.

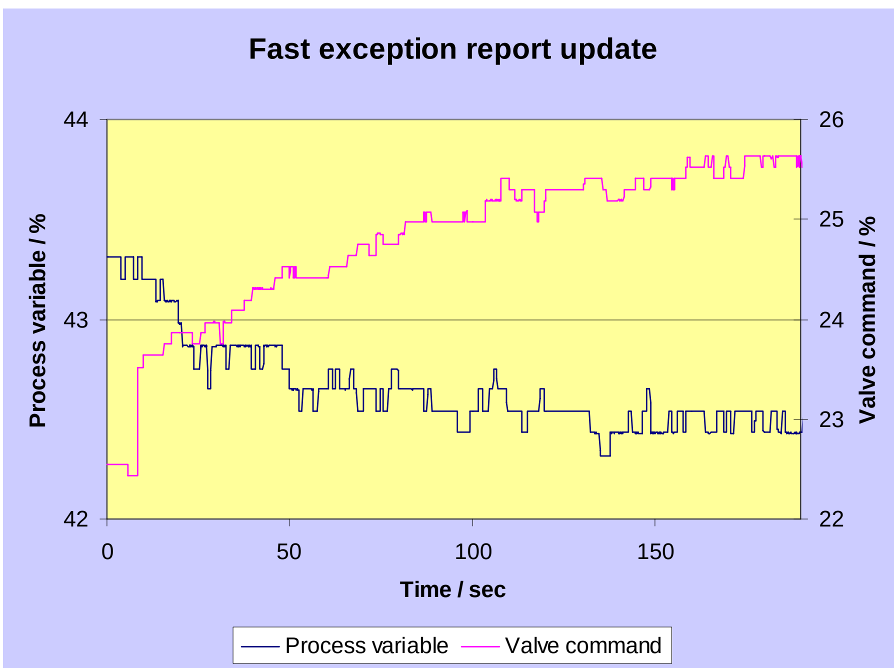

Figure 9. Controller with improved exception reporting

This example illustrates an important point that controllers should not be 'split' across networks unless the controller signals can be updated at high speed. It is also mentioned in passing that the overall control loop integrity is reduced in the 'split' solution by increasing the number of components in the loop, any of which may fail.

Valves

During a recent function test of major control valves on an offshore production platform, 20 valves out of the 35 tested were found to be inadequate. The failures were due to the following reasons:

- 7 valves failed to open until the commands reached 10-20% (5% was considered adequate).
- 9 valves exhibited excessive hysteresis of 10-20% (again 5% was considered adequate).
- 4 valves were 'hunting'.

The performance of 17 of the valves was improved by re-calibration while 3 required more extensive refurbishment.

It should be pointed out that this platform was no worse than most others, in fact arguably better since the operations team had initiated the function test to rectify the valve problems! So if these results can be considered typical, 57% of control valves are not operating correctly, but 85% of the problematic valves can be repaired without major expense.

For comparison, 29 transmitters were also checked on this same platform. The 6 flow, 10 temperature and 8 out of 15 pressure transmitters were working accurately. The calibration of the remaining 5 pressure transmitters was only slightly out. Valves are clearly the weaker link!

Figure 10 illustrates the impact of a faulty control valve. This trend display shows the flow through a compressor and the command from the anti-surge controller to the recycle valve. In simple terms, compressor surge is caused by gas flowing backwards through the compressor and can result in severe damage to the machine. A pipe is installed from the compressor discharge back to the suction and a valve in this line is opened whenever the compressor approaches surge conditions, thus ensuring that gas is available to maintain forward flow and to prevent surge. From the trend display, it can be seen that the gas flow is dropping over a period of several seconds. The controller correctly interprets this as an approach to surge and the recycle valve is commanded to open. Although the command ramps open, the flow continues to decrease. Finally, at 8 seconds, the controller responds to the rapidly approaching surge by progressively stepping open the valve until the flow recovers. Note that the command has reached about 15% before the flow changes direction.

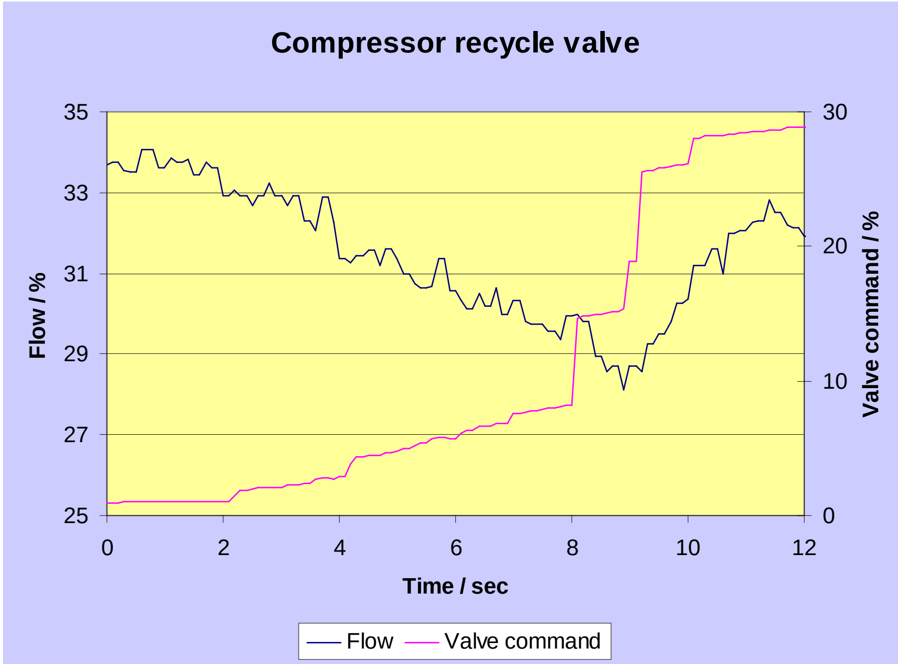

Figure 10. Effects of valve not lifting from seat

The valve is not lifting from its seat until the command reaches about 15% with the following consequences:

- The anti-surge system is failing to provide adequate protection to the compressor. A more rapidly dropping flow may have resulted in surge.

- This particular brand of anti-surge controller is configured to step open the valve in an emergency. Although saving the compressor, this action sends a considerable disturbance through the downstream process and is best avoided.

Figure 11 shows a similar problem on another compressor. Again, the flow is dropping away and the controller responds by attempting to open the valve. No recovery has happened even by the time the command reaches 18%. The controller responds to the approaching surge by stepping open the valve command to 40% giving a large increase in gas flow. Not shown is the subsequent drop in discharge pressure and rise in suction pressure which introduce significant disturbances throughout the plant.

Note that this type of anti-surge controller 'washes out' the recovery step over a period of about 14 seconds. With the flow reduced to its initial level, the system is now primed for a further drop towards surge. Typically, such problems are seen as apparent cycling of the compressor flow due to the recycle valve stepping open every few minutes (again supporting the argument for high speed trends).

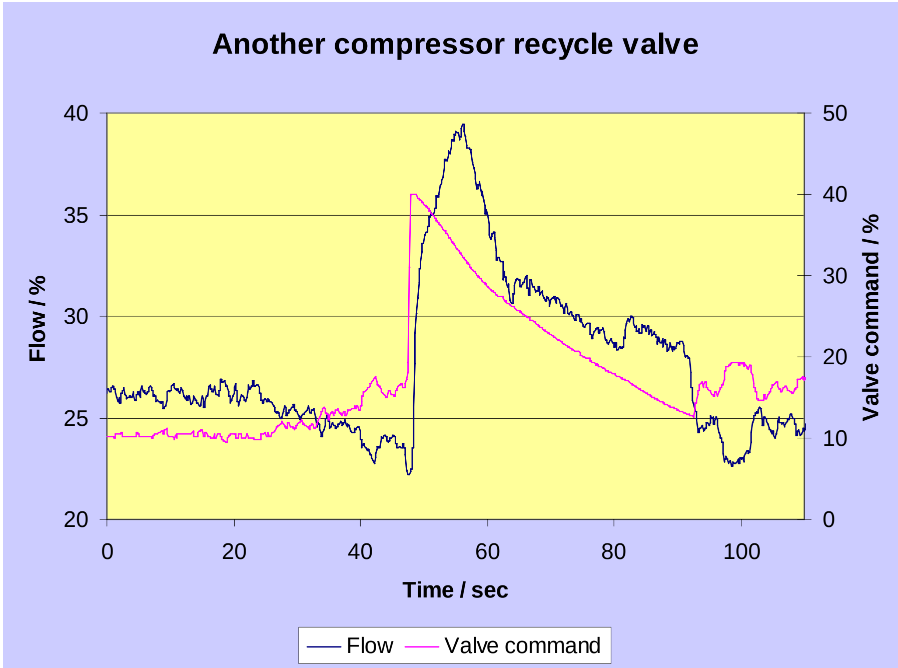

Figure 11. Another valve not lifting smoothly off seat

Figure 12 shows a completely different issue. This is the trend of the flow input and valve command of a flow controller and is exhibiting a cyclic behavior with a period of 25 seconds. Note that the valve command is ramping by only 0.5%, indicating that the valve is experiencing very slight hysteresis. Since it is practically impossible to completely remove hysteresis and stiction effects, the performance of this valve is extremely good. However, this very small hysteresis is resulting in more than 17% change in the flow (almost 4% over the full 0-55 mmscfd range of the flow measurement). The problem here is that the valve is grossly oversized so that insignificant valve movements cause very significant changes to the flow.

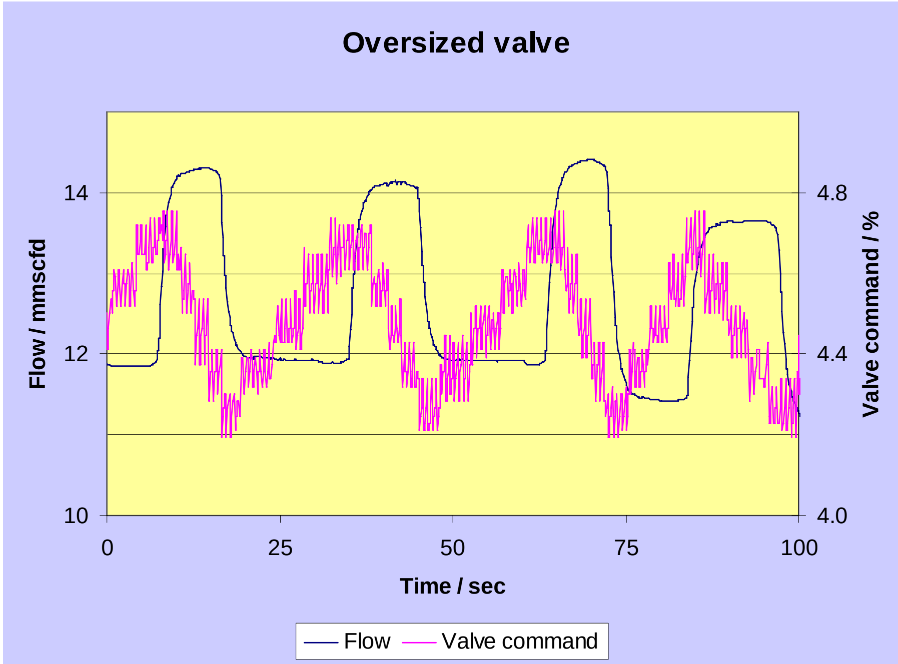

Figure 12. Oversized valve

It is highly unlikely that such small hysteresis can be removed by re-calibrating the positioner and anyway the problem will reoccur in time. The solution is to replace the trim with a more suitable size or consider using a parallel valve arrangement if large flow turndown is required.

Valve Testing

The previous section will have stressed the importance of control valve performance. The normal five-point positional tests are not sufficient to prove that the valve is stroking correctly. The valve must also be tested with several small 1-2% command changes, particularly over the normal operational zone of the valve.

For example, flare and recycle valves are normally closed. Tests should ensure that the valve lifts from its seat and follows small command changes around the fully closed and say 25% open positions. Some inaccuracy at the fully open region will not matter. Similarly, if a level control valve normally operates at about 30-40% open, then test the response of the valve to small command changes around the 35% position.

To ensure a control valve is seating firmly closed, the positioner is often set up to place the valve on its seat at about 4.5 ma of the command's 4-20 ma range. By the time the command has dropped to 4 ma, corresponding to 0%, the valve is fully seated. One common problem is that the output commands of many control systems are configured to wind 5% beyond the 4-20 ma range to ensure valves seat correctly. So with the controller outputting 3.2 ma and the valve not lifting until 4.5 ma, the command will have to increase from -5% to about 3% before the valve begins to respond during which time the process variable will probably have deviated even further from the setpoint.

Summary

Hopefully, this article has convinced you of the necessity to look at the data before leaping in to modify the tuning. Many process instability problems may look superficially like poor tuning, but with correct observation of the data turn out to be due to totally different reasons. It may well be that tuning will alleviate the problem, but the improved insight and knowledge of the real issue will definitely assist in making the correct decisions during the remaining tuning exercise.

Historical trend displays are critical for providing an overview of the controller response and to assist in fault-finding. The trend must supply adequate detail and purchase of suitable high speed data acquisition equipment is strongly recommended.

This article included several examples where controller tuning had been incorrectly blamed for poor control. Of these examples, the valve performance is by far the most common and should always be checked before proceeding. The weir- and network-related problems are rather more unique and were included to illustrate the broad spectrum of potential issues. The main point is to keep an open mind and to observe the data carefully.

Contek Systems Ltd

Contek is an independent process control consultancy located in Aberdeen, UK and has extensive experience in analysing control engineering problems and optimising the performance of controllers in the onshore and offshore oil & gas industries.

Based on a broad practical and theoretical control engineering background, Contek is also the developer of control and mathematical applications for use on Microsoft® Windows.