



31: Designing for Control

Suggested by one of our readers, in the last of his series of articles, Myke King covers examples of lessons learnt from poor decisions made at the process design stage

CONTROL ENGINEERS are frequently called upon to address problems that could have been avoided. Rectifying poor process design is usually impractically costly once it is frozen. Involving an experienced control engineer during the design phase will avoid the problems, often at negligible additional cost. Too late an involvement results in lost opportunities and compromised controller design.

THINK ABOUT THE DYNAMICS

One of the key differences between a control engineer and a process design engineer is the importance placed on process dynamics. Process design engineers are primarily interested in steady state behaviour and would rarely consider the impact their design has on deadtime or lag (*see TCE 981*). For example, Figure 1 shows two possible designs for controlling a simple heat exchanger. From a steady-state perspective, the schemes are interchangeable, and the design engineer would likely choose the lower-cost option requiring the smaller control valve. But, dynamically, the scheme on the right is subject to the thermal inertia of the exchanger and so would require slower controller tuning, resulting in greater temperature variation.

Figure 2 shows a scheme implemented at one of the author's former client sites. The site had a circulating fuel header from which its fired heaters and boilers took their supply. Flow measurements were installed on both supply and return – the idea being that consumption could be determined by the difference between the two. This proved unreliable. For example, fuel consumed on a heater might be around 5% of the flow through the header. So, a metering error of 1% in both flowmeters could therefore result in a 40% error in the calculated difference. Indeed, when fuel consumption was low, the calculated difference could even become negative. Without a reliable measurement, it was simply not possible to install a fuel flow controller.

As discussed in *TCE 1,002/1,003*, there is a dynamic advantage in cascading the temperature controller to a fuel flow controller. Each heater draws fuel from a common header. A large change in fuel consumption by another user disturbs the header pressure and alters the fuel flow to our heater. The resulting effect on heater temperature may not be apparent for several minutes, introducing significant deadtime and lag and requiring slow

QUICK READ

- **Design for dynamics, not just steady state:** Seemingly equivalent designs can have very different controllability once deadtime and lag are considered
- **Instrumentation and architecture choices enable or limit control:** Early decisions determine whether effective cascade, feedforward, inferential, and advanced control can ever be implemented
- **Avoid shortcuts that compromise long-term performance:** “Quick fixes” such as hot gas bypasses or split-ranging often create interaction, inverse response, and tuning limitations that persist for the life of the plant

controller tuning. With no flow controller in place, the change in pressure could not be compensated for immediately, which would have resulted in negligible temperature disturbance.

The installed instrumentation also meant that other beneficial schemes could not be implemented. For example, feedforward control (*see TCE 999*) and excess air control (*see TCE 1,002/1,003*), cannot be installed without a reliable fuel flow measurement.

A commonly installed alternative scheme relies on controlling burner pressure (*see Figure 3*). While giving an early indication of the disturbance, it introduces the problem that it acts in the wrong direction when burners are taken in or out of service.

And don't forget about the dynamics of on-stream analysers. On-stream analysers are frequently included in process design but generally considered only as aids to process monitoring. Including them in a control scheme can be extremely beneficial, provided process dynamics are considered in the process design. This might, for example, influence the choice of analyser. We saw, in *TCE 1001*, the benefit of compensating fuel gas flow for variation in its heating value. While on-stream calorimeters are available, they introduce a measurement delay that may make their inclusion in firing control impractical. Inferring heating value from an in-line densitometer eliminates this delay and is a significantly lower-cost solution. This also makes the installation of multiple analysers practical, for example to enable cross-checking, or to accommodate long fuel gas headers by installing individual analysers closer to each heater.

Sample delay must also be considered. If, because of the location of the analyser house, long sample lines are unavoidable then a fast loop should be installed. For liquids, this is a larger bore pipe that often takes flow from the product pump, returning most of it to the pump suction. It passes close to the analyser house and the analyser sample is then taken from this loop. The delay will then be considerably less than that caused by a lengthy small bore sample line. Using the pressure drop across a control valve to drive the loop should be avoided, as variations in valve position cause changes in loop flow and measurement delay, leading to tuning difficulties. If no suitable pressure drop is available then the fast loop should include a small, dedicated pump that returns to the fast loop off-take.

Taking the sample as far upstream as possible should also be considered. For example, it is common, on distillation columns, to sample the overheads product from the discharge of the product pump. The overhead condenser, reflux drum and associated pipework dramatically impact the process dynamics. A corrective change made to the reflux will have no effect on the measured composition for some considerable time. A controller using the analyser measurement would have to be very slowly tuned. If the column operates with a total condenser, consideration should be given to withdrawing the sample from the column overhead vapour line. It requires a different design approach. Firstly, the sample should be withdrawn from the first elbow via a probe centred in the pipework so that it is representative and not affected by any condensation occurring on the pipe wall. Secondly, the sample line needs to be steam-traced and insulated so that the sample remains in the vapour phase. The sample delay will be then much less than it would be as a liquid. Thirdly, a route needs to be found for the discharge of the fast loop, returning it to part of the process operating at a lower pressure.

AVOID THE DREADED HOT GAS BYPASS

On most distillation columns, condenser duty is manipulated in some way to control pressure. Energy input is primarily determined by the reboiler, the duty of which will be set to deliver the required product compositions. Should the energy removed by the condenser be out of balance then the vapour condensed will not match that produced and the column pressure will change. Controlling pressure is an effective means of maintaining energy balance.

If the condenser uses a liquid coolant, for example water, then duty can be adjusted by varying the flow. For air-fin condensers, manipulating air flow is not so simple. The schemes, shown as Figure 4, include the use of variable speed fans, adjusting the pitch of the fan blades and adjusting louvres. All are costly and usually prove, because of the mechanics involved, unreliable. A common solution is to fix the air flow and install a vapour bypass – known also as a hot gas bypass. Figure 5 shows a commonly installed scheme. Although bypassing the condenser, the vapour is still fully condensed. Opening the bypass causes

Figure 1: Exchanger bypass




Figure 2: Fuel header

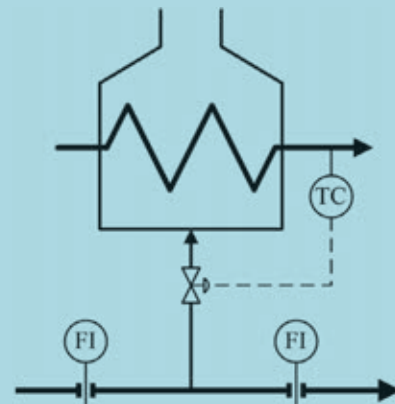


Figure 3: Fuel pressure versus flow control





Figure 4: Air-fin controllers



the vapour to be condensed in the drum where heat transfer is less efficient and so the pressure rises.

From a steady-state perspective the control scheme would appear to be a good solution. However, dynamically, it can present a major tuning problem. Figure 6 shows the result of a simple step test, along the lines of that described in *TCE 981*. As expected, at steady-state, opening the valve increases the pressure but it shows the phenomenon of inverse response. The controlled variable initially moves in the direction opposite to the steady state condition. Installing a PID controller under such circumstances is likely to be impractical. Any correction it makes would result in the controller error initially increasing – causing it to make a yet larger correction. Unless tuned to act very slowly, the controller is likely to be unstable. The problem arises because of the two competing effects of opening the valve. The first is the increase in vapour flow from the column to the drum, reducing the pressure. The second is the reduction in vapour condensation, increasing

the pressure. But, because of the thermal inertia of the condenser, this will be slower. The first effect therefore initially wins out until the resulting increase in drum pressure restores the vapour flow close to its starting point.

Figure 7 shows a partial solution, where opening the valve will cause the drum pressure to increase because of both the increased vapour flow and the reduced condenser duty. It is typical of a “quick fix” that might be considered during commissioning when the problem with the original scheme becomes apparent. While feasible, it is a good example of an enforced compromise. While it clearly holds the drum pressure constant, it permits the more important column pressure to vary.

There are a range of alternative solutions that, because of the additional cost, might be rejected at the process design stage. One is shown as Figure 8. Instead of the valve being placed in the condenser bypass it now directly controls the flow through the condenser. Known as a flooded condenser it is now partially filled with liquid. Opening the valve reduces the liquid level, exposes more tubes and so increases condensation. The bypass line would now be described as a balance line maintaining the drum and column pressure equal.

Perhaps the costliest solution is shown as Figure 9. With both valves installed, both drum and column pressures can be controlled. Rather than rely on the operator to ensure the drum pressure is kept sufficiently below the column pressure, a pressure difference controller is installed. Maintaining this difference ensures that opening the condenser valve increases the condenser duty without increasing the vapour flow to the drum. However, this introduces another problem for the control engineer – one of interaction. Imagine there is some disturbance that increases column pressure. The pressure controller will open the condenser valve to increase condensation. But the pressure difference will also increase. To relieve this, the dynamic process control (dPC) will open the bypass valve – reducing vapour

Figure 5: Hot gas bypass (common scheme)




Figure 6: Inverse response

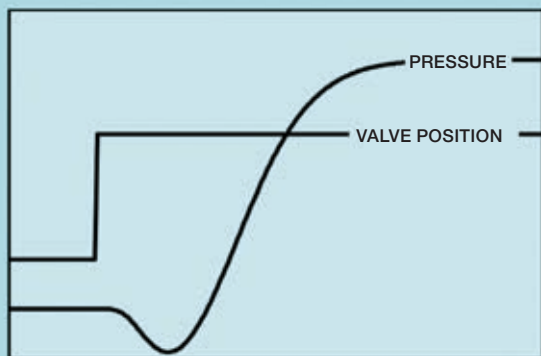


Figure 7: Relocated pressure transmitter





Figure 8: Flooded condenser



condensation and so initially increase the column pressure. The two controllers will fight – requiring one be tuned to act very slowly and so sacrifice stability. There is a “quick fix” whereby the dPC measurement is configured to use the column pressure controller setpoint rather than its measurement. It then only corrects for changes in drum pressure. It breaks the interaction in one direction so that the controllers no longer fight but changes made by the column pressure controller will still affect the dPC.

Of course, the best solution is for the process designer to select a solution other than the hot gas bypass.

Figure 9: Two pressure controllers

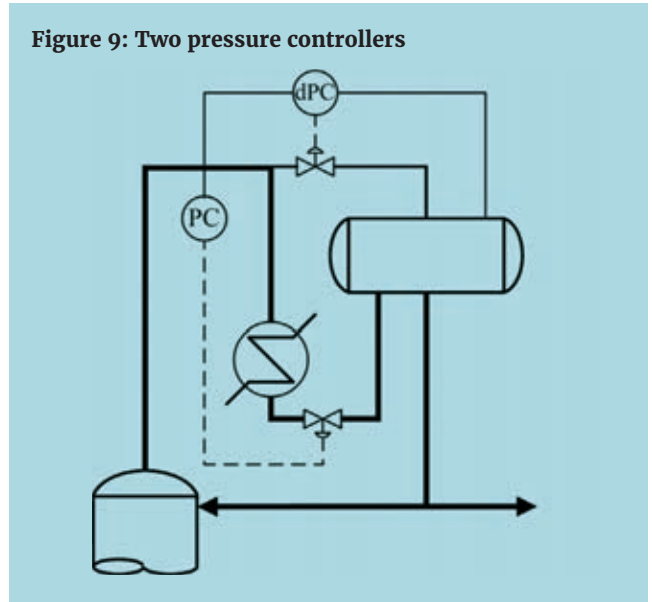
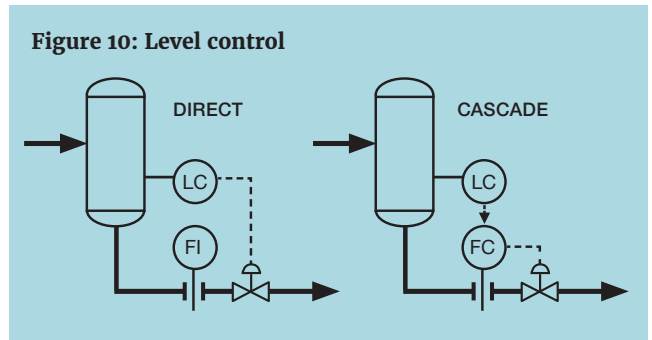


Figure 10: Level control



SURGE CAPACITY

One of the largest improvements that can be made to process stability is the use of surge capacity. This might be deliberately included in the process design – such as a feed surge drum. Or it might be in a vessel that is installed for another purpose – such as the reflux drum we’ll see later in Figure 12. This capacity can be used by tuning the level controller to be averaging rather than tight (see TCEs 987 and 989). However, averaging control usually requires that the level controller be cascaded to a flow controller, as shown in Figure 10. Without this, the controller must be tuned tightly to maintain the level when either the upstream or downstream pressure changes. With the cascade in place, the flow controller will quickly compensate for pressure changes – leaving the averaging level controller to smooth out flow changes.

So, at the process design stage, it is important to decide whether the flow controller is required. Should tight level control be required, for example, on a compressor suction drum then it’s right that direct level control be installed. While it might be useful to have a flow measurement for monitoring purposes, its inclusion in the control scheme will make a critical controller more susceptible to instrument failure.


There are many examples where averaging control would have been advantageous but the process design excluded the necessary flow measurement. Orifice type meters require 30 pipe diameters of straight pipework. This is unlikely to be in place and so retrofitting would be prohibitively costly.

During commissioning, level controllers should be placed in automatic as early as possible. Unlike most other measurements which are self-regulating and naturally reach new steady state, levels are integrating (see TCE 990/991) and so require much more attention when in manual. Most level controllers can be tuned based only on process design data. Such tuning calculations should form part of the process design.

AVOID SPLIT-RANGING

Split-ranging (TCE 1,004) is a technique that extends the operating range of a controller by sequentially adjusting more than one valve. Figure 11 shows the pressure controller of a fuel gas header. Preferentially gas supplied via valve A should be consumed first. If this flow is at maximum, and the valve is 100% open, then valve B is opened as required. In the first scheme, the pressure controller

Figure 11: Split range alternatives



output is split into two ranges. As it changes from 0 to 50%, valve A moves from 0 to 100% open. As the PC output increases from 50 to 100%, valve B opens 0 to 100%. When instrumentation was entirely pneumatic, this scheme had cost advantages in that both valves connected to the same copper pipe. These days, there is no cost advantage but the scheme is still extensively specified in new process designs. It has two disadvantages. The first is that the dynamic response to one valve moving will generally be quite different to the other. Even in this simple example, if the control valves are different sizes or have different upstream pressures, the process gains could be very different. Controller tuning therefore must be a compromise. Secondly it can be difficult to accurately calibrate the valve positioners leading to either an overlap, where both valves move when the controller output changes, or a deadband where neither moves. Again, this can present a tuning challenge and, in some cases, increase operating costs.

Figure 11 shows alternative improvements. The first uses two separate pressure controllers, sharing the same measurement. Each pressure controller can now be optimally tuned. To have them preferentially open valve A, the setpoint of PC1 would set slightly higher than that of PC2. The second scheme employs a valve position controller. It is a conventional PID controller which has the output of the pressure controller as its measurement and the desired position of valve A as its setpoint. If, for example, the current valve position is greater than the setpoint, the VPC will open valve B. This increases the header pressure so that the pressure controller closes valve A until the desired position is reached.

COMPRESSOR SURGE AVOIDANCE

We covered this recently (see TCE 1,013). Many organisations were persuaded to install surge avoidance schemes in a dedicated PLC (programmable logic controller) rather than in the main distributed control system (DCS). While interfaced to the DCS for operator displays, this incurred additional hardware cost and increased maintenance support. Further, most PLC do not offer the preferred version of the PID algorithm. The surge protection technology is

largely in the public domain and can readily be installed in the DCS by the plant owner or by a specialist supplier. Contrary to many compressor manufacturers' claims, it is generally not true that the DCS scan interval is too slow for effective protection.

PLAN FOR INFERENCE PROPERTIES

As we described in TCE 1,005-1,008, inferential properties use measurements of flow, temperature and pressure to predict a key property – usually product quality and most often that from a distillation column. Inferentials can be the first principle type, based on simplified process simulation, or regressed from historical data. An experienced control engineer will identify instrumentation, maybe not otherwise included in the process design, that would greatly improve the accuracy of the inferred property.

Figure 12 shows a typical control design for a binary distillation column. The tray temperature controller is a simple inferential in that the bottoms composition is related to the bubble point measured in the lower section of the column. This relationship is pressure-dependent; to accommodate changes we install a pressure-compensated temperature (PCT), as described in TCE 996. On most columns the pressure measurement used for control can be used in the calculation but this is not the case for vacuum distillation. Here the pressure drop across the column is significant compared to the operating pressure. In addition to that at the top of the column, used for control, a pressure measurement close to the tray temperature should also be installed.

The PCT is largely a measure of cut; it effectively determines how much material leaves the tray as vapour and how much leaves as liquid. But, as we saw in TCE 1,009/1,010, product composition also depends on fractionation. There are measurements, such as reflux flow, which are an indication of fractionation, but their inclusion in the inferential can bring little improvement in accuracy. And, if located well away from the temperature, these additional measurements can also cause difficult dynamics. On most columns, the inclusion of another temperature measurement a few trays from the first can make a substantial improvement to inferential accuracy. It effectively

permits inclusion of the slope of the temperature profile, which relates directly to fractionation.

Adding an additional tray temperature to an operating column is prohibitively costly. Following the welding of the required nozzle, the column's shell must be stress-relieved by raising its temperature to around 650°C – really not practical on an installed column. Including the additional thermowell in the original design brings negligible additional cost. Indeed, it would be valuable to also install two similar measurements close to the top of the column, to help with the inference of the overheads composition.

REMEMBER THE BASICS

While not necessarily part of the process design, there are decisions made by default or during commissioning that can have a long-term negative impact. One example is the choice of control algorithm. We showed (*see TCE 984*) the importance of standardising on the I-PD version of the PID algorithm. Commissioning the plant with the DCS's default algorithm creates a future legacy of a large number of controllers that must be modified and returned. In practice, this may never be completed. The same applies if filtering is unnecessarily included to make trends look smoother, rather than considering if the noise actually presents a problem with the final actuator. Not only does remediation require substantial manpower, it delays the installation of lucrative advanced controls. And, indeed, if such advanced controls are implemented above poorly configured basic controllers the cost of re-engineering them precludes later changing the chosen algorithm.

While temporary tuning constants may be installed during commissioning, systems must be in place to retune controllers properly soon afterwards. A data historian is therefore essential and configured to retain uncompressed data. Software to support model identification and controller tuning, if not already on site, should be included in the project budget.


Controller monitoring should also be in place at commissioning. While there are several packages available, as explained in *TCE 1,016*, ad hoc development of Excel-based tools will likely be more cost-effective.

WHEN TO COMMISSION ADVANCED CONTROLS

These days, advanced control has become almost synonymous with multivariable control. It can be advantageous to include this level of control in the project budget, rather than get it approved separately as a later addition. However, it is a myth that it can be designed as part of the process design. Claims that suitable models can be developed from process simulations are rarely valid, and operating constraints often only become apparent after commissioning. Step-testing will therefore still be required well after startup.

So-called traditional advanced control (sometimes called enhanced control) cover those schemes that can be implemented at the DCS level. These would include, for example, feedforward

Figure 12: Typical distillation column controls



control, deadtime compensation, nonlinear control, overrides etc. As part of the process design, they should be specified in sufficient detail to ensure sufficient DCS capacity. Implementation will usually follow process performance testing, but the required manpower should be included in the project budget to ensure the necessary expertise remains available.

TRAINING

Finally, don't stint on the training. On an operating unit, the cost of properly training a control engineer is likely to be recovered within a few hours of the engineer implementing a recently learned technique. On a unit being designed, avoiding poor decisions will ensure that the process can be optimally controlled without any costly retrofits. ■

*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*