Advanced Control: The Next Challenge

Presented by Myke King, Whitehouse Consulting

Whitehouse

Whitehouse Consulting
The White House
Main Road
Dibden
Hampshire
SO45 5TB
United Kingdom

tel	+44 (0)23 8084 0677
fax	+44 (0)23 8084 0673
e-mail	mykeking@compuserve.com

Introduction

If you have started to read this paper hoping to find some extraordinary insight into the next generation of advanced control technology then I am regret that you are going to be disappointed. While of course there will always be technological advances, the main challenge for the oil industry is not in assimilating the latest offering. The vast majority of the industry has yet to fully come to terms with what is already available, indeed many sites have yet to properly apply technologies that have been with us for 30 years! This paper instead aims to highlight what many seem to overlook when applying advanced control, to try and explain why and to offer some ideas as to how the industry can better benefit from its application.

Basic Controls

A good control engineer knows that the basic controls need to operate properly before advanced control can be successfully installed. The industry is generally, although not universally, good at ensuring that the controllers are *mechanically* sound. Certainly controllers that simply cannot be used are usually dealt with quickly. The main problems arise more from using the wrong control algorithm and from implementing the wrong tuning. So why do these problems persist?

Firstly there is not always awareness that the DCS actually offers a wide range of control algorithms. Or, if the engineer has read the manual thoroughly, he may not appreciate why there are so many choices and which should be used under what circumstances. Not fully appreciating the benefit of each algorithm the engineer will select the default or the one that most closely matches his understanding of conventional PID control. In almost every circumstance this will lead to the wrong selection. The engineer will have missed the opportunity to significantly improve the response of the process to disturbances - needlessly extending, typically by a factor of three, the time that the unit takes to recover. Consider subjecting the process in Figure 1 to a reduction in feed rate caused by the operator reducing the set point of FC1.

Figure 1 Process flow diagram

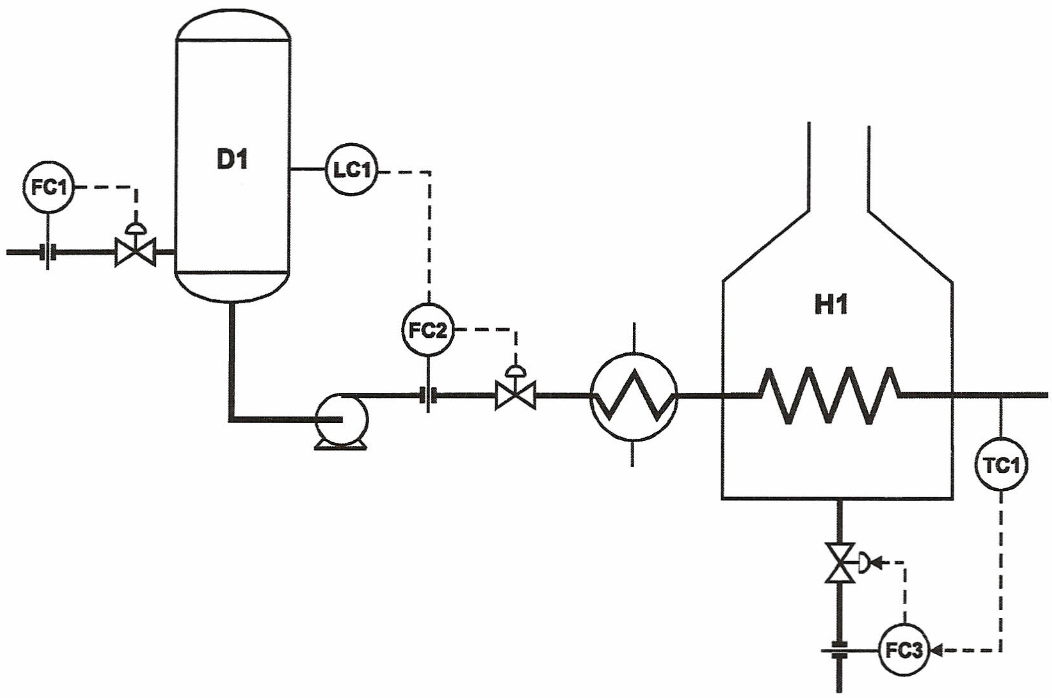

The two heater outlet temperature trends (algorithm A and B in Figure 2) are in response to the same disturbance to feed rate. In both cases the chosen algorithm was tuned to give virtually the same response to temperature set-point changes. Why anybody would choose algorithm A .... but that is exactly the one that most will implement!

Figure 2 Choosing the correct algorithm

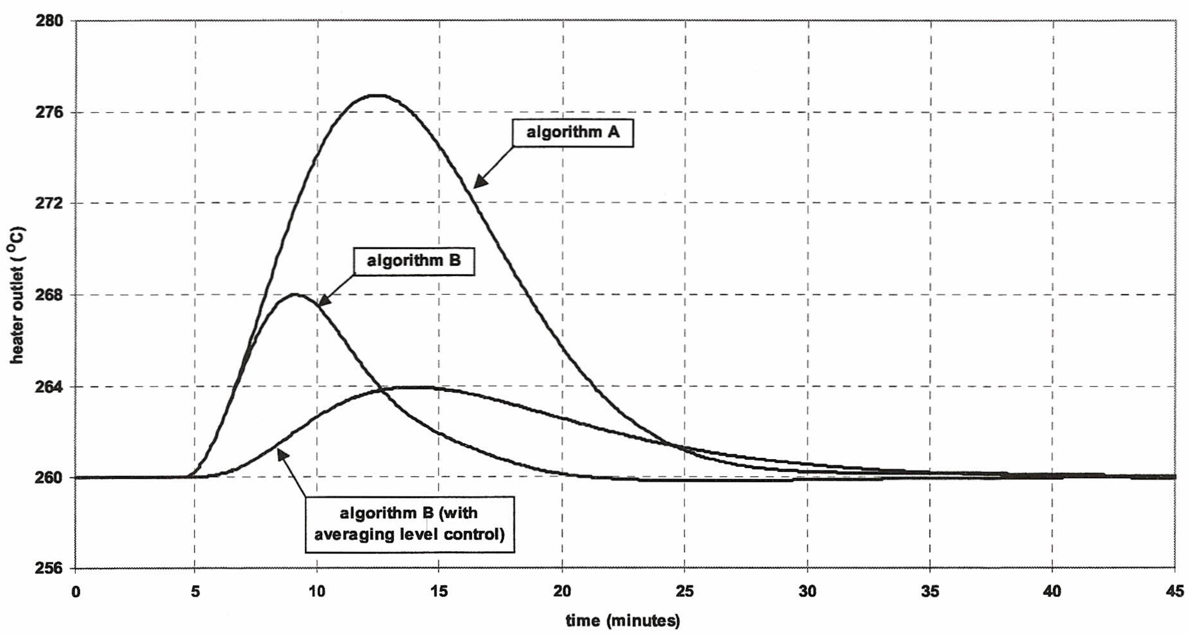

The next challenge is ensuring the basic controllers are properly tuned. Many controllers are not. Probably the most common example is "averaging" level control. This permits the level to vary (between alarm limits), thus minimising changes to the downstream flow. There are many situations where level controllers can be used to exploit surge capacity within the process. This results in much reduced downstream flow disturbances - often resulting in remarkable improvements to process stability. Many control engineers appreciate this, but few properly calculate or properly implement the correct tuning. Further, few truly understand how to exploit the adaptive nature of the "error squared" algorithm. Figure 2 also shows the effect of changing from tight to averaging control.

Step testing, to support multivariable control design, is a major undertaking. It can take several weeks, working shifts to cover round-the-clock testing. It is not something that one would willingly undertake more often than absolutely necessary. Switching basic controller algorithms, or re-tuning them, changes the process dynamics. Once step testing has started the engineer has effectively committed the refinery to retaining poor basic controllers at least until the next major process modification, when step testing would have to be repeated in any case.

So how do we prevent this situation from arising? The first priority is the right training. The DCS vendors are generally poor at explaining why their systems support so many different algorithms. Indeed it is likely that the majority of vendor engineers do not understand the reasons themselves. Further the APC vendors are often less demanding than they should be in terms of the performance of the basic controls. They will check that the controls are "about right" but will not wish to become involved in any significantly time-consuming modifications. After algorithm selection the engineer is next faced with a bewildering array of tuning methods. Almost every month a journal somewhere publishes a new one or a new product is announced. With only a few exceptions these are usually flawed. Rarely do they account for the variety of algorithm types and usually they apply incorrect tuning criteria. The industry needs a definitive tuning aid.

Finally we must make sure that the engineer has the time to implement the necessary changes. There is no doubt that properly setting up basic controllers can be very time-consuming, maybe even tedious. We need to ensure the engineer focuses on the controllers that justify attention. Implementing an optimum design that halves the time taken by a flow controller to deal with a process disturbance is of little value if the controller already deals with it in a matter of seconds! We also need to ensure that the project schedule includes sufficient time. This will usually mean not even committing to the implementation of advanced control until the basic work has been completed. APC vendors are under their own internal schedule and cost control pressures. They will not want to delay step testing - least of all, repeat any.

Inferential Properties

Also known as “soft sensors” or “virtual analysers”, inferential property calculations are an important part of any advanced control project. While the technology is often applied to predicting product qualities from more easily measurable operating conditions, it can be applied to almost any variable that is difficult or costly to measure directly. Examples include column flooding, coker drum outage, catalyst activity etc. The technologies available range from simple linear formulae, through neural networks to semi-rigorous engineering calculations.

There are numerous suppliers of the technology. Some offer highly functional (and highly costly!) design software that, by analysing historically collected laboratory and process data, generates the “best” correlation. Others argue strongly that a far better approach is to apply process engineering knowledge and develop the calculation from first principles. There are arguments for and against both approaches. There is certainly a place for both. What the refinery overlooks is the accuracy of the calculation. Most engineers will “eyeball” line graphs and scatter charts, to see how the predicted property compares to that measured. Some will even look at statistical parameters such Pearson R^2. But few will examine whether the accuracy is sufficient to capture the benefits. Even fewer will continue to monitor it after commissioning to ensure that accuracy does not degrade. And perhaps the most frequent error is to implement automatic updating by laboratory result in the mistaken belief that this will improve accuracy; it will usually make it worse! Figure 3 illustrates this. When initially commissioned the controller worked correctly in driving the process until the inferential reached the HI limit. However, due to the inaccuracy in the inferential the operator noted that laboratory results were frequently exceeding the HI limit and so reduced it to compensate;

reducing to the point where the quality give-away was similar to that before commissioning! The controller designer responded to this problem by implemented automatic laboratory updating. This introduced larger errors in the inferential ultimately resulting in the operator reducing the HI limit yet further. Laboratory updating is the equivalent of assuming a company's share price is the same as yesterday's - generally a reasonably accurate estimate but not of much help in predicting sharp falls.

Figure 3 Inferential Property Control

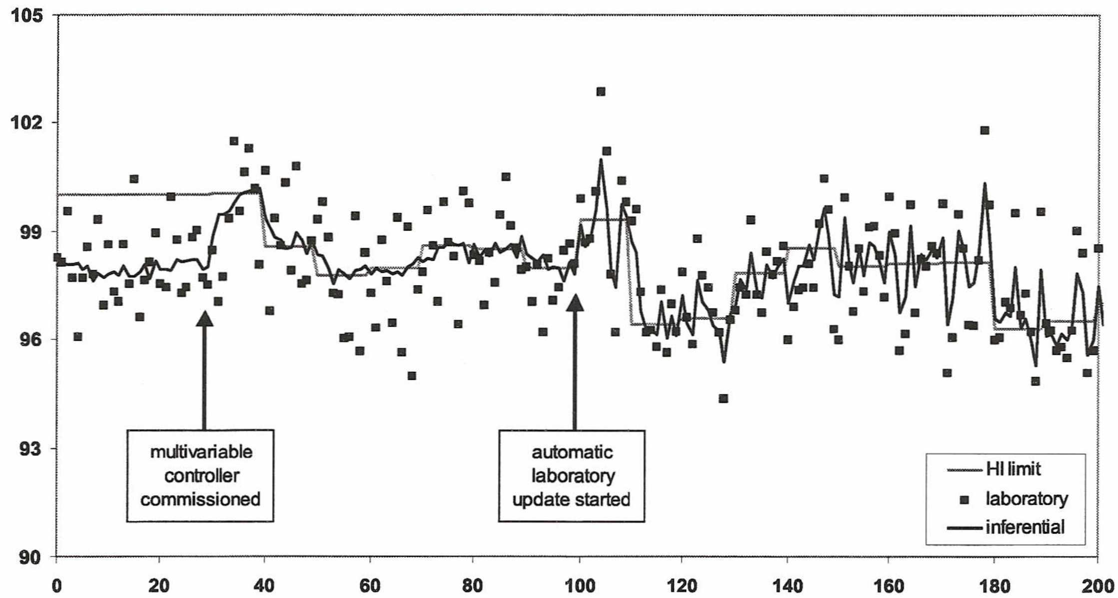

Such problems usually arise because the inferential calculation has been derived or calibrated using historically collected process data. Such data may be unreliable, particularly if laboratory data are not accurately time-stamped or the process not at steady state. Often there is not enough variation in process conditions to enable the effect of each measurement to be separated from this "noise". It may be that the data were collected during plant tests; the resulting model may have fitted the data collected at the time but does not work well with the current operation.

Refiners are often seduced by costly software packages that have impressive user interfaces and a large number of features. While the cost of developing these may justify the high license costs, too often these add little to the *value* of the package. A well engineered inferential, developed in a spreadsheet or similar product, can outperform one that has been produced by naively applying a more costly package.

Usually the best approach is to begin to develop inferentials several months before the start of the advanced control project proper. While the *elapsed* development time is likely to be lengthy, the work involved is relatively little. With a little guidance most refineries have the resources to do the work themselves or they can bring in one of the specialist suppliers. Most APC vendors will incorporate inferentials developed by others. If for some reason they are reluctant to do so, then at least the refinery has set a benchmark for the vendor to improve upon using his own technology.

If an inferential technique cannot be applied then there may be no alternative but to install additional instrumentation. The lead time for such work is usually many months. Not identifying its need until during the step testing means that the controller will be commissioned without it. No one will wish to delay the project until the instrument work is complete. Indeed it may be that the work is *never* done, with the resultant sacrifice of benefits. At best it will be necessary to repeat much of the step testing once the instrumentation is available. This is an additional argument for exploring the feasibility of inferential techniques well in advance of the project proper, in order to maximize the likelihood that the necessary instrumentation will have been identified and installed in time for the step tests.

Multivariable control design pitfalls

Despite the technology being in common use now for a couple of decades there are still a large number of controllers still being implemented with potentially serious design deficiencies.

The controller may later drive the process away from the conditions under which the step tests were performed and the step testing may not have taken account of any non-linearities in the process. The changes in process gains may be sufficient for the controller to drive to the wrong constraints. Further, controllers may not to adhere to the rules of mass balance if process gains are implemented in the controller without a detailed check. The result is that the controller "believes" it can increase the total mass of products out of the same mass of feed! If it moves towards what is perceived to be an incorrect operating strategy, the engineer will adjust the cost coefficients or maybe delete sub-models until the controller behaves as required. While this might capture the benefits, the approach is risky. In future, the controller could well drive the process *away* from optimum operation.

This brings us to another important issue. The controller should if possible use *real* process economics. It should include all feed and product flows as either MV's or CV's. Any other variable that can cause changes in utility consumption and have a substantial effect on the overall profit function might also be included. The cost coefficients should be applied primarily to these variables and be recognizable feed costs, product prices, utility costs etc. The controller will then make *objective* decisions. It is not bound by years of tradition. Provided the process gains are correct and complete then the chances are that the strategy it implements is correct, even if this is different from the established one.

There are several examples of spectacular increases in profitability that have arisen from placing more faith in the controller. Unfortunately there are far more examples of engineers adjusting the cost coefficients to force the controller to follow the expected strategy and therefore reducing process profitability *below* that before controller commissioning. Figure 4 illustrates this.

Figure 4 Distillation Column Optimisation

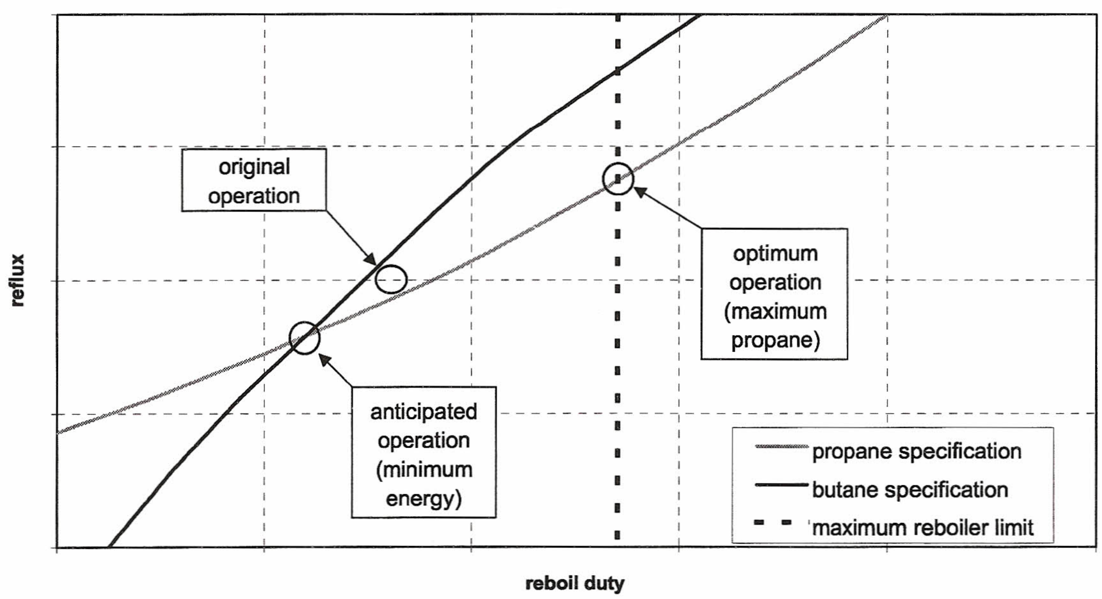

On a propane/butane splitter it was anticipated that the controller would drive both products to their purity specifications and hence improve unit profitability by reducing energy consumption. However commissioning the controller with real product prices and reboil cost caused it to drive towards maximum propane recovery, of course consuming more energy but with this additional cost more than compensated for by the increased production. Rather than check whether the controller had identified a more economic operation, the commissioning engineer artificially increased the reboiler energy cost until the controller followed the established strategy, thus better achieving a loss-making objective!

It is important that the refinery's planning and economics group is closely involved in the APC design and commissioning. The group will naturally tend to consider advanced control outside of their understanding and responsibilities. While the dynamic behaviour is not relevant, the steady state behaviour certainly is. The controller is, after all, a LP but more importantly, an LP that directly manipulates the process. The group needs to understand its general principles and be actively involved in setting many of the objective coefficients.

The APC vendor will not usually consider it part of his responsibilities to challenge existing operating strategies. It is important that refinery personnel closely involve themselves with the project to ensure that all key economic variables are included and to validate the process gains and economics used in the controller.

Most engineers are aware of the need to regularly check the constraints towards which the controller is driving. Process operators tend to reduce the range in which MV's are constrained. This is usually in response to some problem perceived with the moves that the controller is making. Left unchecked, the majority of MV's become fixed at limits and the controller has effectively been taken off control. Simply trending the number of constraining MV's will help detect the problem. Adding economic information greatly enhances the effectiveness of monitoring. Recent releases of controllers now allow reduced costs and shadow prices to be monitored. The use of real economics in the controller then permits the cost of artificially imposed constraints to be determined. Faced with quantified economic arguments, Operations Department are less likely to impose costly conservative limits. Should the constraint prove to be realistic then the same economic calculations provide an instant estimate of the incentive to debottleneck the process.

Retaining management support

The are two fundamental challenges in retaining the correct level of expertise. The first is convincing refinery management of the need to do so long after project completion. There exists a mistaken belief that, once a multivariable controller is commissioned, the ongoing support requires substantially less expertise than its implementation. While there may on occasions be a case for reducing the *number* of personnel it is unlikely that the *quality* of expertise can be reduced without jeopardizing long term benefit capture. Ensuring that the controller still properly models process behaviour, that it using the correct economics and that it is not being over-constrained requires not only continuous support, by high quality support. If a controller indicates that a more profitable operating strategy should be adopted; strong entrepreneurial skills are needed to challenge long held beliefs to the contrary. Further, even relatively minor process changes can require almost complete re-engineering of the controller - demanding the same level of expertise as the original project.

The more senior the manager championing advanced control investment, the more successful is the implementation likely to be. The more senior this manager,

- the greater his authority in approving investment or the greater his influence on those that do;
- the more easily can the advanced control project manager persuade the organisation to release the best engineers to the project team;
- and the more attention will the advanced control vendor give to winning the project and to the quality of its execution.

Sooner or later most refiners reach the stage where a senior manager is appointed who has a strong belief in the importance of advanced control and the authority to bring about its implementation. Such a manager will ask the question “Why aren’t we investing in APC?” rather than “Why should we invest in APC?” What follows is usually a fast-track project, often having a noticeable effect on the “bottom line”. However, in many organisations, the duration of the project can outlast that of the manager. A site-wide project can take several years to approve and execute, particularly if its scope also includes OM&S and process information systems.

It is not uncommon for a project to be completed under a different manager to that which initiated it. The problem can arise where the successor is not so aware of the value of advanced control and/or of the need to retain high quality support. He is likely to allow the level of expertise to decline. He will be less inclined to approve the funds necessary to keep the technology current, let alone extend its use to other units or implement higher levels of optimisation. While he too may only hold his position for a few years, during his tenure the benefits captured will fall as controllers are less able to be effective as process operation and economics change. Even if his successor is another champion of advanced control, it will take the organisation several years to again build up the expertise and to update the technology that has fallen into disuse. By the time the refinery is ready to move to higher levels of optimisation, this manager could have also moved on! Figure 5 illustrates the cyclical benefit capture that arises from this "stop-go" approach.

Figure 5 Impact of Management Change

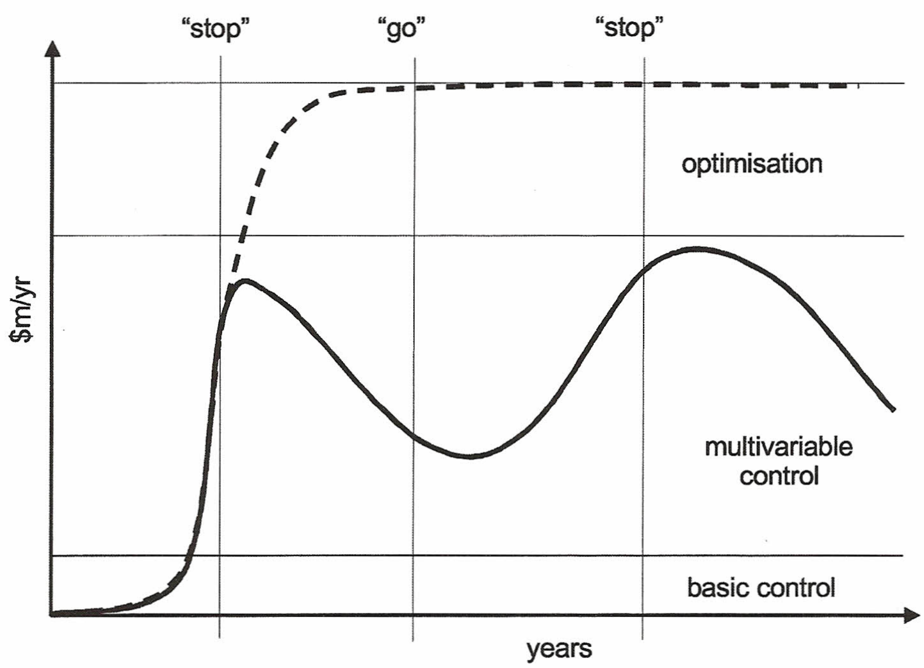

Anything that can be done to shorten the project will therefore be beneficial; completing it under the same manager will enable the organisation to quickly move towards higher levels of optimisation. Committing to site-wide implementation of the full range of technologies should help maintain the momentum through management changes.

Organisations can become desensitised to the benefits captured by advanced control. A successful advanced control project, even without a detailed post-project benefit analysis, will make *noticeable* changes to process profitability in a very short period – usually within days of commissioning. Once achieved benefits become less obvious. Switching off a controller does not result in an immediate equivalent loss of benefits. The gradual degradation of control performance, caused by neglect, will be even less noticeable. Indeed this desensitisation can cause a self-fulfilling prophecy. Because the organisation is not aware of the benefits it is less inclined to assign the right level of support resource, which ultimately causes a real loss of benefit. It is important therefore to maintain a high profile of the success of advanced control. Key is routine monitoring, regular inter-department reviews of advanced control performance and frequent publicity (in a management digestible form) of successes or failures. Comparison with other sites can also help.

Retaining expertise

The second challenge is deciding how best to provide the expertise. For decades now refiners have been assigning high quality process engineers to advanced control projects with the intent of using the project as a vehicle to train the engineer in advanced control technologies. While this aim is usually met, the refiner then finds that the engineers enjoy the work, become frustrated that the refinery has no immediate projects of a similar nature and leave to join the APC vendor. Some refiners respond to this by assigning lower quality staff who are unlikely to be offered employment by the vendor; this again jeopardizes long term benefit capture. Alternatively they rely more heavily on the vendor to provide the support. While, depending on the quality of vendor personnel assigned, this may enable the project to be successfully implemented it does not provide the long term support necessary. There is no guarantee that the same level of expertise can be provided by the vendor. Indeed, in recent years the flow of good engineers appears to have reversed in favour of the refiners. Many engineers find the demand to work away from home has increased to the point where working for a refiner is preferable. Coupled with the falling workload being experienced by some of the leading vendors, now is a good time for refiners to recruit experienced engineers.

With good people management practices, refineries could improve the retention of process control expertise. Within many companies more kudos is attached to management careers than to technical careers. Those staff climbing the technical promotion ladder are too often perceived as failed managers. While it is unlikely that the industry will ever offer salaries to technical staff commensurate with those of their management peers, much can be done to encourage high quality personnel to stay with the company. Key to this is rotation through the process control group.

It should be clear to others that joining this group does not mean "career death". Indeed the opposite should apply. The best process control engineers well understand refinery operation and economics, and will have strong interpersonal skills. If not wishing to follow a technical career, they should ultimately become senior managers. Indeed there is an argument that a pre-condition of becoming such a manager is a proven understanding of the contribution that process control can make to refinery profitability.

Rotating personnel through other departments is beneficial. Personnel from Operations Department will bring a practical perspective to the group and, after a year or two, can return to operations in an influential role as champions of the technology. Those assigned from groups other than Process Engineering (generally the best source of trainee control engineers) can surprise the organisation, and themselves, and become effective long-term control engineers.

Summary

Around 80% of the major process units in Europe have some form of advanced control. Assuming typical advanced control benefits of 20 cents/barrel, European refining capacity should be capturing around $800k/yr. The reality is that many of these controllers are capturing little or, worse, driving the unit *further away* from optimum operation. In reality only about half of the opportunities are actually being exploited. On the basis that there are around 100 refineries, in an average refinery there are $4m/yr of benefits available from simply making better use of the technology that is already installed.