Multivariable Control

Hans H. Eder
ACT

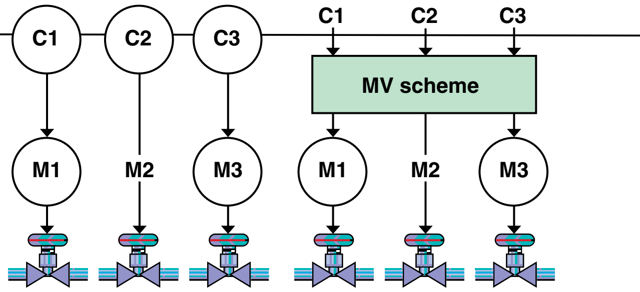

FIGURE 1. SDecising whether to use single loops (left) or a MV scheme (right) depends, in part, on how stronly the variable are coupled

DECK DECK

Multivariable (MV) control is a powerful technology. Applied in industry for several decades now, substantial benefits from its use have been reported in many publications. Yet, this technique is still by far not utilized to the extent it could and should be, one reason being is that there is still a lot of uncertainty about it: Many control professionals are not proficient enough on what MV control really means and when to use it. In fact, most consider it as a rather complex, expensive and time-consuming endeavour, altogether inherently too heavy to be dealt with – a perception that calls for correction.

Therefore, the motivation of this article is to make it more clear what multivariable control is all about, to position it and consequently to help it to a better use for the benefit of the industry. Discussed will be where and when MV control should be used, how to select the proper approach and what must be specifically observed to ensure a successful application. And to make this all a bit more tangible, two real life applications will be presented as well.

The basics

Multivariable control is a technique that allows us to deal with more than one control objective at the same time. For a particular piece of equipment or a process unit two or more variables, so-called controlled variables (Cs), must be kept at their target values, their setpoints, and to be able to achieve this several handles to interact with the process, manipulated variables (Ms), are needed. It does not matter here whether M is a controller (for instance, a flow controller) or a final control element, (such as a valve); important is that there are at least as many Ms present as there are Cs.

As a first attempt we could consider mastering this system just with individual, isolated loops where variable C1 is controlled by adjusting M1, C2 by acting upon M2, and so on. This of course will work if there are no effects from one loop's action upon the other. However, if M2 not only influences C2 but also C1, then any intended change in M2 will cause an unwanted change in C1. Loop 1 will bring C1 eventually back to its setpoint but in doing that most likely now C2 will be effected. This ongoing mutual disturbing reduces the performance of the loops, it increases the variation of the controlled variables around their target values and can lead, in the worst case, to total instability. So, clearly a different solution is needed.

Multivariable control provides the mechanisms that allow us to change the setpoint of one C without disturbing the other ones, that is, without driving them away from their setpoints. This effect is called decoupling. We could therefore also say that MV control is a technique that provides the decoupling of interacting variables. To achieve this, the control actions cannot be computed individually any more, as in the case where there is only one single input and one single output of the controller, the SISO case. Instead, the actions of all manipulated variables must be determined in a coordinated way, in one global control scheme that considers all the Cs and Ms, a scheme with multiple inputs and multiple outputs (MIMO).

Figure 1 compares a situation where there are three controlled variables and three manipulated variables. On the left side, the Cs are controlled by individual loops, where M1 and M3 are slave controllers of cascades, while the controller for C2 is directly outputting to a valve. On the right side of Figure 1, MV control is applied. The measured values and setpoints for all controlled variables C1 through C3 are inputs into the MV controller whose outputs are the required values for all the Ms — the setpoints of M1 and M3 and the valve position of M2.

When to use MV control

From the above example, we can deduct that MV control is useful and necessary when:

- More than one control objective needs to be met, and
- Interactions exist between the process variables involved

To make this a bit more practical, let us consider a simple and seemingly trivial example, the blending of two material streams. Both streams consist of pure components, namely A and B. The task is to deliver a final product in the required quantity (flowrate) and also the required quality (that is, with a certain concentration of component A). The control objectives in this case are: 1. The total flow rate, which we will call the controlled variable C1; and 2. The concentration of component A in the final product, our controlled variable C2.

There are two variables that we can manipulate in order to meet these two objectives, namely the flowrate of component A, here defined as the manipulated variable M1, and the flow rate of component B, the manipulated variable M2. Now let us see how this system behaves.

When the flowrate of component A (M1) changes, then the total flow C1 will change, but the concentration C2 will also change. Likewise, upon every

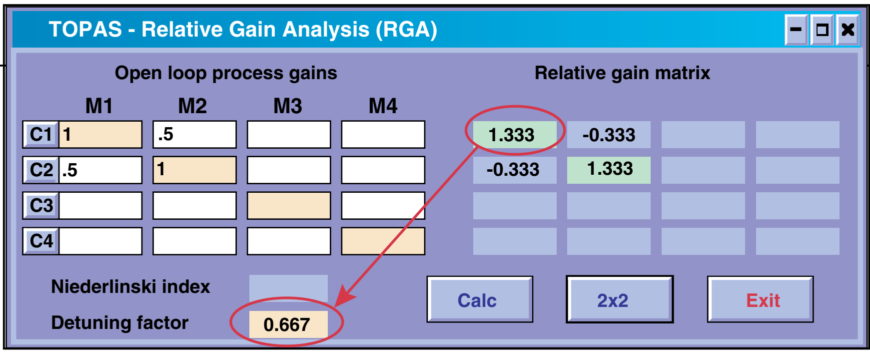

FIGURE 2. Relative gain analysis and the Niederlinski Index are the quantitative factors determining the degree of interactions between variables

change in M2 not only C2 but also C1 is effected; there are clearly unwanted effects present, "interactions". In other words, when one manipulated variable is changed to correct for a deviation in one control objective this causes unavoidable and unwanted change in the second one.

Because of this situation, we must decide if we intend to meet the two control objectives by two individual control loops or by means of a MV scheme. For this, we need to answer two questions:

1. How strong are these interactions?
2. How should the controlled and the manipulated variables be paired? Is it useful to control C1 with M1 and C2 with M2 or better vice versa?

Regarding the first question we could simply say: When the interactions are "strong enough" then we need to "take care" of them. But this is a very qualitative, fuzzy criteria. Much better of course is to base the decision on quantitative facts.

Such facts are delivered by a method called RGA, the relative gain analysis (the term I prefer) or relative gain array. To explain this method, we use the blending example from above and assume that there are two loops provided to control C1 and C2 by manipulating M1 and M2. We want to investigate loop 1 (C1 – M1) and compare one test where both loops are open with another one where only the second loop (C2 – M2) is closed. For both tests we determine the effect of the step change in M1 on C1, the process gain. The relative gain is defined as the ratio of these two process gains. This is the decision variable we are looking for.

If the process gain is the same for both tests, then the relative gain is 1. That means that there are no interactions between the variables and they can be controlled by individual loops. The further the relative gain departs from 1, the stronger are the interactions. Typically, for relative-gain values greater than 1.5 and lower than 0.5, we will not use single, isolated loops any more.

It should be noted that, even in cases where the interactions are relatively weak and individual loops are used, the PID controllers need to be detuned for stability reasons (compared to the case with no interactions at all, the basis for all PID tuning calculations). This means a reduction in performance, which can be avoided with MV control.

A second indicator, the Niederlinski Index, NI, helps to check if the pairing of the variables (the decision to control C1 with M1 and C2 with M2) is in order. We will not go into the details here, but important to know is: Whenever the NI is negative, then the tested pairing will be inherently instable – and thus useless.

In our example it turns out that we would need to pair C1, the total flowrate, best with component A as manipulated variable M1 and thus C2 with M2 for cases where the target concentration of A in the product is low, and the other way round for high concentration targets. This, however, leads to different control schemes depending on the target property of the product, which is not a desirable situation at all — one more motivation for using MV control.

Figure 2 shows the RGA results for a two-by-two case (2 inputs and 2 outputs) and the required PID detuning. Although the relative gain is just 1.3, the proportional gain of the PID controllers must be lowered by one third.

A simple approach: IAC

Once the decision for MV control has been made, the next question to be addressed is: How can these interactions be compensated? How could these loops be "decoupled"?

In principle, it is quite simple. With respect to variable C1, M2 is nothing else but a disturbance, acting upon the process output. To compensate this effect, we have a well know means, namely the feedforward. Thus, in order to keep variable C1 always at its setpoint we need to provide a feedforward that "reads" the change in M2 and makes adequate adjustments in M1. The same holds true for the situation of C2. This leads to a construction where the two feedback loops are now seconded by two crosswise arranged feedforwards. This control scheme is called inter-acting control (IAC).

For the most simple case with just two Cs and two Ms the resulting scheme is pretty simple, yet there is a serious draw back, namely the tuning of the application. Because of the ever present deviations in the feedforward parameters from reality, the disturbances can never be fully compensated and thus go "round and round." We do not want to get deeper, just state the final verdict: While a PID-based IAC scheme for a 2 x 2 case is feasible, simple and built in no time, IAC is not recommended because the effort needed to reach reasonable performance is not justifiable. As a consequence of this experience, MV control leads us inevitable into the domain of model-based control. Let us look now at this technique.

IAC versus model-based control

Figure 3 shows a process with interactions controlled by an IAC scheme. On the right side, the process is represented by the dynamic-transfer functions that describe the effect of M1 upon C1 (G11) and that of M2 on C2 (G22) plus the interactions, that is, the effect of M2 on C1 (G12) and the effect of M1 on C2 (G21). The IAC system is shown on the left side: For each controlled variable (C1 and C2) there is a PID controller, their control functionality being described by the functions C11 and C22 respectively; and there are two feedforwards, each taking the output of one PID as input, computing the compensation action according to transfer functions C12 and C21 and adding this action to the other PID's output.

In the case of a model-based control-

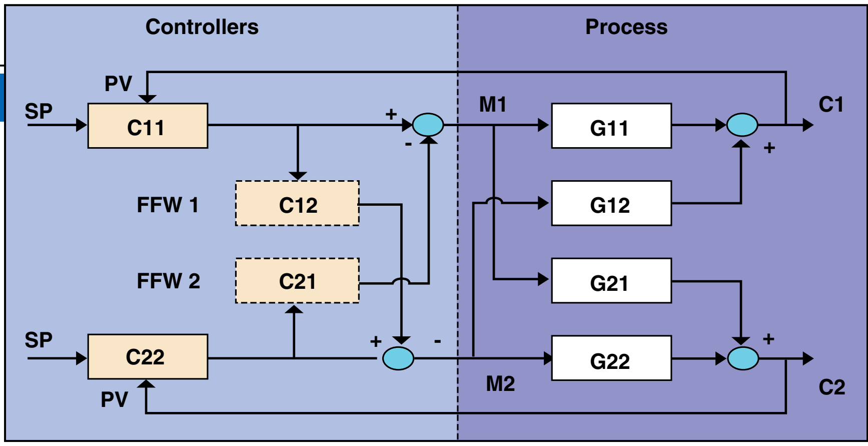

FIGURE 3. The scheme shown here is for interacting control (IAC), one of the simpler methods for decouipling loops

ler, we are not using a PID algorithm to calculate the changes in the controller output, but rather a suitable representation of the process behavior, typically in the form of equations or parameters. These models can be formulated in quite different ways and thus the first, most fundamental decision concerns the type of model to be used.

Model types

Process models range from those based on first principles to true black-box formulations. Models of the first type are built on sound knowledge but are typically much more complex than those using the second approach.

Black-box models contain no knowledge about the fundamental workings of the process, they just describe relationships between variables. If we have two time series of data, for instance one for the heat input into a process and one for the resulting temperature, it is under certain circumstances conceivable that the best match is achieved with a model that says that increasing the heat will decrease the temperature. Of course, this makes no sense at all. Therefore we never should rely solely on statistics, but always apply our process know-how and judgement. The result we call "grey box models".

The main questions to be answered for the decision regarding the model type are:

1. What is the targeted range for the application? Is it relatively narrow and well known or an operating region never explored before. Black-box models can only be trusted within the range used for their development while a white-box model describes the process beyond the familiar boundaries
2. How much will the conditions in the process change over time? Black-box models deliver a snapshot of the process. They can be made adaptive, in some cases even with little effort, but do not describe the inherent mechanisms that lead to the changes in the deadtime, time constants or process gain. By contrast, white-box models do.

We have thus to decide between fidelity and flexibility on one side and simplicity on the other. In most cases, and especially for larger MV systems, preference is given to the simpler approach, the grey-box formulations.

Also, the formulation of the model and especially the description of the process dynamics can be done in quite different ways and several different approaches, from most simple to most complex, exist to compute the sought result, the values of the Ms. In the following, we will look at two examples, one from either end of the spectrum.

We start with the most simple case, a 2x2 scheme using a linear, static model. We can formulate the relationship between the Ms and Cs simply by writing

$$C1 = a1 + a11 * M1 + a12 * M2 \quad (1)$$

$$C2 = a2 + a21 * M1 + a22 * M2 \quad (2)$$

where the parameters a11 represents the process gain of the effect of M1 on C1, a12 the gain for the effect M2 on C1 and so on, and a1 and a2 are bias terms.

The equations represent the mathematical problem to solve. In process-control terms it means we have to compute the values of M1 and M2 that bring and keep both C1 and C2 at their target values. Of course there is more to MV control than solving such a trivial set of equations. In the above example there are no dynamics in the model (these are taken care of elsewhere), there is no constraint handling foreseen, nothing provided for proper initialization, no means to tune the controller, to make it faster or slower. But we want to show here just the bare principle.

When we have just 2 or 3 variables we can easily compute the solution. When there are more then we need the help of special tools, equation solvers. In the majority of the cases we are using linear models and therefore we are in the world of linear programming (LP), and the most widely used algorithm for solving sets of simultaneous linear equations is the so-called Simplex algorithm. Thus let us take a quick look at this elegant and most efficient tool, that uses very little resources and can be used practically in any DCS.

LP formulation

First we re-arrange Equations 1 and 2, moving the bias terms a11 and a21 to the left hand side. Now only the Ms and their process gains are on the right hand side.

$$C1(setpoint) - a11 = a12 * M1 + a12 * M2 \quad (3)$$

$$C2(setpoint) - a21 = a22 * M1 + a22 * M2 \quad (4)$$

The typical LP formulation is shown in figure 4. Its core is the matrix that contains all the process gains. For every M there is one column and for ever C one row. Additional "artificial" variables can and will be used and thus we normally have more rows and columns present.

To the left there is a vector that holds what is shown in Equations 3 and 4 on the left hand side. Funny enough, it is called the Right Hand Side (RHS) vector. Then we have two vectors UBV and LBV that hold the maximum and the minimum allowable values for the Ms and below them the solution vector will give us the needed answer, the values of the Ms that satisfy the Cs setpoints. On the top we have one more vector where we can put penalties on the use of the variables in the columns. This is the way to drive the solution into the desired direction when extra degrees of freedom exist.

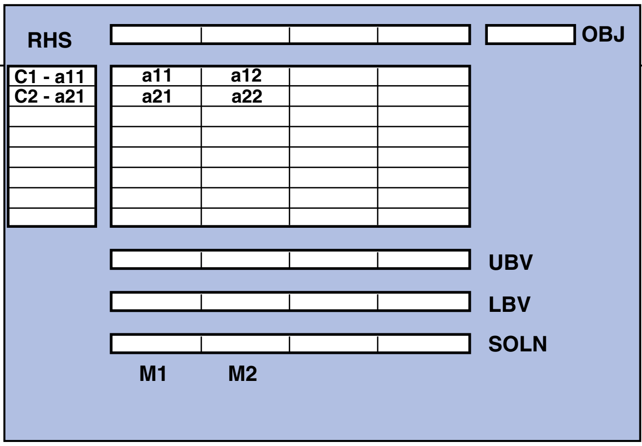

FIGURE 4. A typical linear programming (LP) scheme

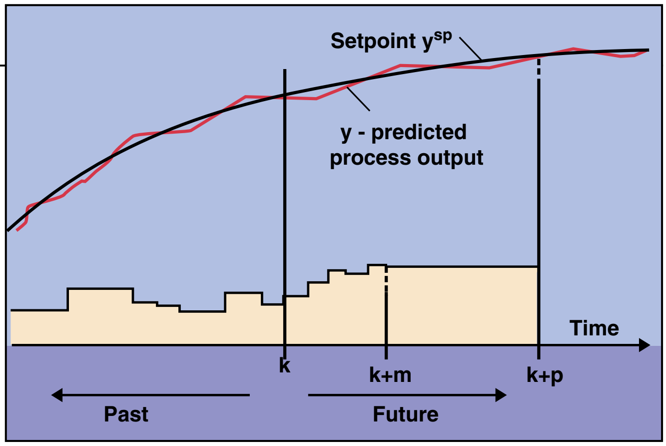

FIGURE 5. The dynamic matrix control (DMC) approach

In other words, the LP has inherently optimization capabilities.

Without going further into detail we can state that just one control scheme can handle many different situations — just by adjusting parameters — and can be used for regulatory control as well as constraint control as well as for optimization. This opens new possibilities, because the LP allows to handle different operating scenarios for which normally different control applications are required within one single control scheme. Besides, every LP-based control scheme uses the same structure and thus can be easily understood by any new (LP-proficient) user. This approach for solving the MV-control task represents the most simple case.

On the other side of the spectrum is he DMC (dynamic matrix control) approach. In DMC, every relationship is described by a series of parameters k_i, the step response at every time point. Summing up these parameters over time delivers the current value of the controlled variable (= the process output) in question, assuming that a step response has settled after n samples

$$y(t) = y_0 + \sum_{i=1}^{n} k_i * \Delta u(t-i) \qquad (5)$$

where k_i are the samples of the process unit step response and Δu is the change in the process input.

As shown in Figure 5, a trajectory can be defined by the user which the process is supposed to follow towards the new setpoint and the control action is computed by finding a compromise in minimizing the deviations from the trajectory on one side and avoiding overactive movements of the manipulated variable(s) on the other. This is achieved by the objective function

$$Obj = \min \left\{ \begin{array}{l} \sum_{i=1}^{p} \gamma_i^2 [y^{SP}(k+i) - y(k+i)]^2 \\ + \sum_{j=1}^{m-1} \lambda_j^2 \; [\Delta u(k+j)]^2 \end{array} \right\} \qquad (6)$$

where p is the total prediction horizon and m is the control horizon. The terms γ_i and λ_j are weighting factors that allow the tuning of the controller.

The chosen model formulation allows to describe any conceivable shape of a response curve. However, the system needs now to manipulate some 40 matrices in order to compute the solution. Thus, this approach cannot be used in the DCS any more.

Furthermore, in the LP case described above, manual improvement of the model by the user is quite simple and effective. In the DMC case (and also, for example, for Neural Net formulations) this is not possible: When the process has changed and the model is not adequate any more, then it has to be re-identified entirely. More flexibility has its price.

Some commercial systems

Just to round off the subject some commercial algorithms and systems shall be mentioned here. One of the first commercial MV system was HIECON, marketed by Adersa, a French company founded by J. Richalet, who's known as the "grand father" of model based control. In the early 1980s, DMC appeared, developed by Cutler and Ramaker and marketed later by Cutler through his company DMC Corp. In 1996, DMC was bought by Aspen Technology, a company specialized in simulations. Today AspenTech markets DMCPlus as their main package. Several other vendors to mention are Honeywell, which offer a product called RMPCT, Invensis with Connoisseur, Shell offering SMOC, and Emerson, the only one to provide with DeltaV Predict such a solution in the DCS.

All the mentioned system are relatively big and (with the exception of Emerson) require an extra computing platform. They are thus not in line with an old postulate of process control that is strongly supported by the author: The existence of a complete Technology Set with algos and/or schemes for model-based control (SISO and MV), model-free optimization, constraint control and so on, right in the DCS itself. As an alternative to these big systems he (WHO?) developed AMC, to our knowledge the only model-based control system that can run practically in any DCS.

Typical MV control application

MV-control schemes need to gather lots of information and be able to conduct extra checks before and after computing the changes to the process. All these tasks and functions have to be properly organized, best in several distinct parts.

Part 1. Information gathering: Just as any controller, MV systems need as inputs all the current process variables (PVs), setpoints, clamps and other restrictions, but also the operating modes of all involved variables and lower level controllers. As said earlier, MV schemes offer great flexibility and thus typically must deal with many different scenarios. For this, extra information has to be provided, for example which of the Ms are available and which not, ans so on.

Part 2. Problem setup: Based on the information gathered in Part 1, now the problem formulation can be performed and the relevant parameters set. This has nothing to do with adaptive control — we are not changing the tuning of the controller, but we adjust the formulation to match the current situation in the plant. Also in this Part, the actual driving forces are set depending on the momentary objectives and scenario, for example regulatory control or optimization, in the latter case maximization of minimization of the objective function, and so on.

Part 3. Calculation of the solution: The problem formulation is then handed over to the equation solver whose task it is to find a feasible and unique solution. A remark needs to be made here concerning the two words "feasible" and "unique." First, the solution (the values of all the Ms) must be within the given boundaries. Since this is not always possible, the equations must be formulated in a special way to allow the LP (or other solver) to "buy" extra room to manoeuvre (for example deviations from targets) in order to avoid impossible (in LP jargon, "infeasible") answers. Furthermore, we must avoid "alternate solutions", that means an infinitive number of possibilities to satisfy the stated objectives. The consequence would be that the process is permanently moved without any improvement in the targets.

Although LPs are used in many fields, such as operations planning, this task requires special skills because, contrary to all other disciplines, in closed-loop process control, there is no human involved in analyzing the solution and, where necessary, correcting and re-submitting the problem into the solver. This all has to work automatically and ensuring usable solutions is a real challenge for the control engineer. By the way, it is good practice to take statistics of how often the solution had to be discarded. This tells us how well the problem is set up — or not.

Part 4. Analysis and implementation of the solution: The next task is to check if the solver has actually delivered a usable solution and also that the situation in the plant is still the same as fed into the solver. We have seen huge systems where it took hours to compute the solution. During this time many changes can occur in the plant making the solution obsolete. This aspect is a strong motivation to keep the application as compact as possible.

Now comes an important point. When the solution has been found to be "good", it must be implemented in full, else it must be discarded. Partial implementation of the solution of any MV problem is an absolute taboo because all the calculated changes in the Ms are needed to meet all the targets and objectives.

With this we know how the application is structured and what is needed to make it function in closed loop. Let us now take a look at such applications in real life and what else has to be considered besides the proper formulation.

Practical application examples

I would like to describe here two applications that tell us about the chances but also about some pitfalls in applying this technology.

Example 1. A fluid catalytic cracking unit (FCCU): This example shall specifically demonstrate the flexibility of the MV, as well as the problems associated with a too ambitious approach. A fluid catalytic cracker is used to break heavy, long hydrocarbon molecules — heavy fuel oil, which is not in great demand — into smaller ones to be sold as gas, gasoline, domestic heating oil and so on. The heart of the FCCU are two big vessels: A reactor, where the feed is brought in contact with the catalyst; and a regenerator, where the coke on the catalyst is burnt off that is formed in the cracking process.

To control this unit, three key objectives have to be met: The reactor temperature, the regenerator temperature and the composition of the off-gas leaving the regenerator. They can be influenced by manipulating the feed rate, the feed temperature, the catalyst circulation rate, the air to the regenerator and an auxiliary heat source called torch oil. This gives in the base case 3 control objectives and 5 manipulated variables and thus 2 extra degrees of freedom that could be used, for example, for optimization.

Controlling three variables seems not too tough a job. Yet, in addition to the interactions there is another complication, the ever changing situation — the feed rate can be variable or fixed, some feed produces little coke and the feed preheat furnace is fired hard plus torch-oil injected, but with other feedstock the furnace is shut down and no torch-oil is used. Besides, the catalyst-circulation rate could be temporarily fixed out of other reasons.

The technical challenge here is to cope with the ever changing scenarios, especially the ever changing number of available manipulated variables, ranging from 2 to 5. Faced with this, already back in 1979 we choose an LP-based control scheme because it allows one to adapt the control scheme for the situation by simply setting a few parameters.

The application worked fine from the beginning, but we had overlooked one key factor, the human operator. As typical for complex units, every shift had developed its own strategy to control the FCCU. Besides, operators would always just make one change at a time. The MV scheme changed now up to 5 setpoints simultaneously and of course it moved them in a way that was neither easy understand or explain, nor did it ever match exactly the strategy of the present shift. Thus the real challenge was to achieve the acceptance of the operator. We succeeded by proving step by step that every possible single pairing of the variables worked correctly – and consequently the entire scheme. A long and cumbersome way. Of course, we would have reached our goal much faster if we would have started with a different MV application first, such as the one described below.

Example 2. A distillation (splitter) tower: This example represents an ideal pilot application for MV control. The task of a splitter tower is to separate the feed into two product streams. The standard basic controls are:

- Pressure control
- Level control for the accumulator drum and the tower bottom
- Control of suitably located temperatures in the rectification section and the stripping section.

The operations objectives are: 1. To

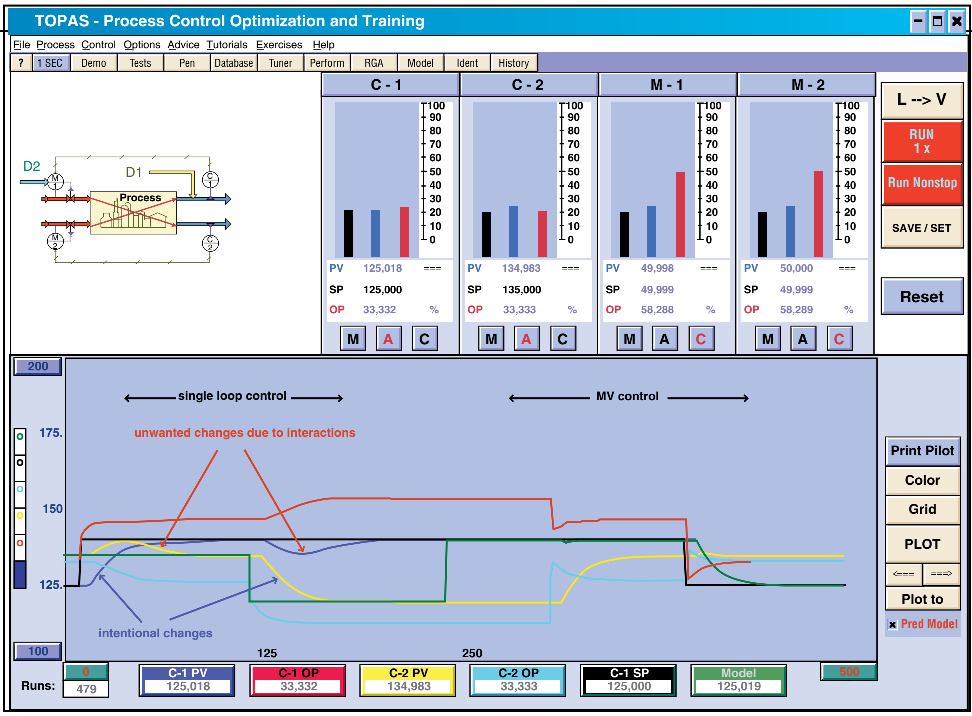

FIGURE 6. Performance of single loops (left) versus MV control (right)

ensure stable operation and 2. To deliver the product streams at specification and with minimum quality giveaway. Let us assume that the product properties are provided on line by stream analyzers and that the temperature in the rectification section is controlled by manipulating the reflux flow and the temperature in the stripping section by manipulating the heat input into the tower.

In addressing the first task (stable operations) all above mentioned controllers are involved and the controller type used is practically always the PID. For adjusting the product properties, the setpoints of the two temperature controllers are adjusted by the operator based on the analyzer readings. Closed-loop control of the product properties is usually not provided in the base case, one reason being that the long inherent deadtimes cannot be handled properly by the standard PID.

Interactions between the temperature loops are clearly present. When the upper temperature is changed to adjust the quality of the top stream, this has some effect on the lower part of the tower and consequently on the bottom stream properties and vice versa. As a result, frequent adjustments in the temperatures are needed and the variations in the product properties lead to some quality give away – and consequently economical loss.

This situation can be improved by means of MV control, one widely found solution being that the control scheme has the two product properties as inputs (control objectives) and the temperature controller setpoints as outputs. Of course, feedforwards for the feed rate, and so on, etc. can be added if needed.

This application is quite simple and small, can be developed with little effort and done right in the DCS, thus avoiding the need for extra computing platforms and interfaces. It is easy to understand, therefore well accepted by the operators and can bring remarkable incentives — all facts that make it a perfect starter. With MV control, there will be less variation in the product qualities, less give away and consequently a more profitable operation — a success that can be also relatively easy measured and communicated.

To demonstrate the performance advantages of a simple MV scheme a screen shot from the tool TOPAS is given in Figure 6. We see a 2x2 system that is controlled in the base case by two separates cascades. Whenever one setpoint changes the second PV deviates for some time from its target. The two lasts tests are done with the AMC controller active and, apart from small effects coming from model inaccuracies, the decoupling is achieved and thus the variations in the PVs are significantly reduced.

Getting started

Opportunities for MV control abound in most industries and plants. Yet, they may not be easily recognizable by newcomers. We need to look for them, but in doing so we never should apply a “technology push”, that is, looking for a place where the technology in question could be force-fitted. The preferred way is to start with questions about the plant’s most pressing needs, it's business drivers and the economical impact of changes: What is the worth of a 1 % increase in capacity or in the yield of the prime product; how much is a reduction of 1 % in fuel, steam or catalyst consumption worth? The answers show where the biggest economical leverage is and thus point us at areas where we can be quite sure that also management will appreciate any positive contributions. After all, they have to agree to the effort and cost involved.

Ideal would be to conduct a short plant-wide study with the objective to identify all improvement opportunities and to define the needed control applications to exploit them. The results are compiled in a document that describes each application in very few words: Their objectives, inputs, outputs, the strategy, and the technology to be used and so on. It will serve later as basis for the detailed development but also tells us what technology shall be used and when. It is the basis for planning our future work.

Now, with several MV controls on our list, work could start. However, when model-based or MV control has never been done in that plant, then the selecting of the very first application is an important and non-trivial task. On one side it needs to clearly demonstrate substantial and measurable improvement, but on the other hand must not take endless time and not overwhelm the developers and especially not the operators.

Sometimes the most attractive opportunities are found where a larger application scope is needed, for example, covering a whole train of distillation towers. Although there is in principle nothing wrong in controlling and optimizing the entire train, it is not advisable to do this as the very first application in that plant. Even if the technical staff could handle it, it will certainly not find spontaneous acceptance of the operators – as the example with FCC application has shown. The punishment for overlooking this barrier will be an unacceptably low service factor.

Some more practical aspects

Besides the selection criteria for the very first application, there are some key aspects and also pitfalls to be observed in the case of MV control that can easily escape the attention of a newcomer. Of course, we just can mention here just a few.

Simplicity – a prime postulate: Always, also when there is sufficient experience with MV control we should strive to keep the complexity of the application low – just as Albert Einstein demanded: "All models should be as simple as possible – but no simpler". The message is clear; we must include the absolutely needed variables, parameters and functionality, but should definitely refrain from "gold plated" designs.

In fact, the most common cause for failure in MV control is over-complexity of the application! We have seen cases where it was attempted to model an entire steam cracker based on first principles with some 500 and 600 thousand differential equations. Even the most powerful plant computer at that time would have taken about four hours to solve them all – too long to expect that the solution is still applicable.

Vendor selection: With few exceptions, manufacturing companies do not have their own MV technology. Thus they have to use outside products and often also outside expertise. Therefore, selecting the most appropriate vendor is as crucial as selecting the right pilot application. To be able to make a sound decision some prior knowledge about the various approaches and their practical implications is needed and taking a truly practice oriented training course on model based control beforehand is advisable. Several key criteria need to be observed:

- The technology: Flexibility, ease of use, robustness against model errors are key factors, but, again, the choice should also depend on the level of in-house expertise. Another factor is the maintenance effort; some approaches allow quite simple adaptation of the model, other require full re-identification
- The tools: Tools for process analysis and parameter estimation are an important factor concerning the quality of the models and also the effort and time needed. In some cases third party tools are superior than those coming with the MV system
- The project approach: Offers range from tight integration of the end-users (strongly recommended) to full turn key projects, the latter ones with the danger that at the end "the key is turned and the customer is locked out"
- The achieved success: Vendors should not only be asked for achievements in terms of variance reduction, but specifically for operator acceptance, for typical application service factors, that is, the percentage of the time the applications are in use. The service factor is a key, perhaps the best success indicator

Advisory mode: Sometimes we are faced with the request to operate an application in the beginning such that the results are first inspected (and eventually adapted) by the operators before passing them to the lower level controls. The idea behind is to make the operators feel more confident before running the application in fully automatic mode. However, this type of operation is not recommendable at all, it has a serious catch; as said earlier, for MV systems the solution must be either implement in full or discarded. It is absolutely impossible just to pick some values from the solution and ignore the rest, let alone to make individual adjustments. The results are inherently poor if not catastrophic and can lead in the worst case to the abortion of the application.

Final remarks

MV control is a powerful technique that can bring significant incentives, many of them stemming from reducing variations in process variables or product properties that are caused by interactions between the variables. Key success factors are the selection of the right technology and vendor, a convincing first application and avoiding over-complex schemes and formulations. Furthermore, if a chance exists we should aim the build the application right in the DCS.

Although in principle very small applications could be done with PID controllers, in practice MV control is done based on dynamic process models. And as for any model based control application, these controls deliver also non-negligible beneficial side effects: The development of the model forces an in-depth investigation of the process and operations. In doing this, always valuable insight is gained and as a result sometimes even strategies or targets are changed drastically. Because of that, we have experienced cases where the whole effort had already paid out before the control scheme was put into full operation.

Furthermore, with such a process model available, off-line studies can be done when unusual situations or operations types need to be explored. All this are strong motivators for MV control, despite the fact that it requires extra training, technology, effort and often also computing platforms. This all means extra expenses and cost upfront, but they are typically paid back in a few months by the application - when it is done right. Just as for advanced process control (APC) in general, we also can say for multivariable control (provided that an adequate, not overly complex MV system applied) that it is the most cost and time effective way to improve the profitability of the plant. ■

Edited by Gerald Ondrey

Author

Hans Eder is president of ACT (Madeloefjeslaan 13, B-3080 Tervuren, Belgium. Phone: +32-2-767-0895; Email: ACTGmbH@compuserve.com), a company that specializes in supplying services and products for process automation and optimization, mostly to major international firms. Prior to founding ACT, Eder spent 21 years with various U.S. and European operations of Exxon, most recently as advanced control manager and computer-integrated-manufacturing (CIM) advisor. Eder has developed software tools, written many journal articles, is coauthor of the Dutch language handbook on process automation, "Handbook Procesautomatie," published by Wolters Kluwer (Amsterdam), anmd has chaired or otherwisde participated in numerous committees and user groups related to process control. He holds a master's degree in mechanical engineering from the Technical University of Vienna (Austria).