WHITE PAPER

# An Enhanced Iterative Process for Maintaining APC Applications

*Michael Harmse, Product Management, Aspen Technology, Inc.*
*Quinn Zheng, Research and Development, Aspen Technology, Inc.*
*Robert Golightly, Industry Marketing, Aspen Technology, Inc.*

# Table of Contents

# Background

A well-engineered Aspen DMCplus® application generally delivers substantial financial benefits via feed rate increases, better conversion, higher yield of valuable products, and reduced energy consumption. However, processes often change over time, requiring modification of the models. This is very problematic, because the cost of modifying the model can be roughly 80% of the initial cost.

# One Plus One Equals Three

When a mismatch between the model and actual plant performance exists, we need to revise the model to eliminate the mismatch.

The workflow is as follows:

- Collect current and representative data
- Pre-process the data to remove bad values and other data anomalies
- Produce a set of candidate models and evaluate the accuracy of those models with due consideration of dynamics, gain ratios and relative gains (RGA numbers)
- Adequately test the candidate models considering the current constraint sets
- Configure and deploy the controller

Most major commercial APC vendors offer tools to assist in these tasks. These tools have provided good value by reducing time and effort, but the potential benefits that can be achieved by making these tools work more synergistically are far greater than the benefits delivered by the individual tools.

# A Conflict of Interest

There is little doubt that automated testing tools deliver good data and reduce the need for continuous monitoring by engineering. In spite of these benefits, step testing is still seen by many as a necessary evil. One issue is the price of generating accurate data. The current generation of automated testing gives us a binary choice: operate the plant with the controller turned on to accrue benefits or turn off the controller to collect open-loop data—and give up the benefits during the test. Because of the limitations of model identification software, it is necessary to turn off the controller to prevent current input values from being correlated with past output values through feedback.

With the controller turned off, the amount of product giveaway could potentially make the test very expensive. This creates a conflict of interest: plant operators and managers want the test over with as soon as possible while engineering wants to ensure that the generation of data produces high information content. How do we deal with this conflict? How can we generate data suitable for modeling while continuing to optimize the operation of the plant?

# Co-Existent Test and Control

If we need a controller (model) to achieve some level of process optimization during the test, that implies the need to safely step test a process unit with an inaccurate model. Can we start with low accuracy models, missing model curves and establish an iterative process of re-identifying the model matrix? That drives the need for a controller to tolerate large amounts of model mismatch.

It's known that when the mismatch grows large enough, the steady-state targets calculated by the linear program may start jumping around. Weak process handles may suddenly be used, leading to large movement in these handles. This occurs when the controller exploits degrees of freedom that require large MV movements for a small improvement in the benefits.

A standard MPC application typically calculates a new LP steady-state target every single cycle, no matter how insignificant the incremental changes are in the cost function. The new steady-state targets are then used to calculate a new dynamic move plan, which guarantees that the controller has integration action and can reject disturbances and remove offset caused by gain errors. This can sometimes cause an overreaction to measurement noise, or lead to large MV movement if near co-linearity is present in the model matrix (i.e., the RGA repair work has not been done).

Tuning parameters are available within Aspen DMCplus to deal with these issues. Moderate accuracy models can be used if sensible MV move suppression values are used. In addition, controller robustness can be traded off with LP optimization performance by using a lower RGA repair threshold in the offline SmartAudit tool. The use of the LP Economic Relaxation parameter and MV/CV cost ranks can help to prevent the LP from flipping from one solution to the next, for small incremental benefits. In addition, prediction error filtering can be used to prevent overreaction to noise and cycles present in the CVs. As a last resort, the selective use of MV minimum move variables can always stabilize a controller with substantial model mismatch.

# Background Testing

It has been observed that an Aspen DMCplus controller in the presence of a very large model mismatch will behave in a much more stable way when all MVs are set to minimum move instead of minimum cost, like in SmartStep multi-test mode. However, when the controller stops optimizing, it only cares about constraint protection. This indicates that pushing for full LP optimization every cycle that exploits ALL the degrees of freedom may not always produce the best result.

There should be a middle ground between running the controller with minimum cost setting and running the minimum move setting, so that one can specify the degree of trade-off between optimization and robustness. In particular, this middle ground approach opens the door to the use of background, low-impact step tests for periodic model updating purposes. The challenge is determining when the engine should run in full optimization mode (every MV on minimum cost setting) and when it can run in constraint control mode (minimum move MVs). To make the data usable for model identification, we would like to impose small amplitude test signals onto the normal control action, but in such a way that the CV deviation will be in the safe direction (away from the active constraint). The tricky part is deciding when and how large the MV step can be such that the temporary relaxation in the optimization result (economic giveaway) is constrained to an acceptable range.

# Closed-Loop Subspace Identification

It has been shown that there are significant differences in the efficiency and accuracy of the various identification algorithms. The choice can have a dramatic impact on the amount of model ID data required. There are also significant advantages for the algorithms that can accurately capture process dynamics and/or deal with co-linearity. The table below illustrates the characteristics of a subset of available identification technologies.

| | MIMO? | MISO? | Non-Linear Optimaztion Required? | Catch Fast Dynamics? | Catch Slow Dynamics? |
|---|---|---|---|---|---|
| FIR | No | Yes | No | No | No |
| Prediction Error (PEM) | No | Yes | Yes | No | No |
| ARX | No | Yes | No | Yes | No |
| Subspace (SS ID) | Yes | Yes | Yes | Yes | Yes |

There are also differences in the algorithms with respect to the use of closed-loop data. The graphs below illustrate the difference in Subspace vs. FIR. We can leverage the advantages of Subspace identification to utilize background step tests for periodic model updating purposes. The graphs below contrast the differences in model quality between Subspace and FIR identification when some degree of closed-loop data is utilized.

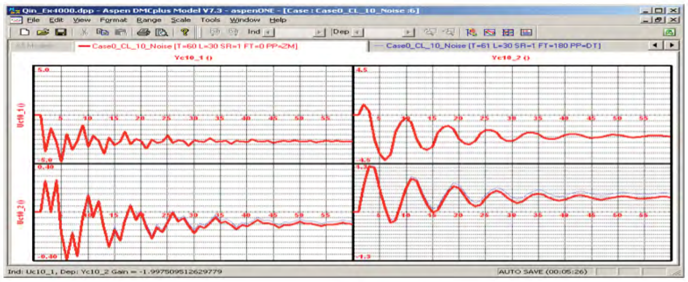

*Identification with PRBS-Test. Subspace ID: Closed-Loop with 10% Noise.*

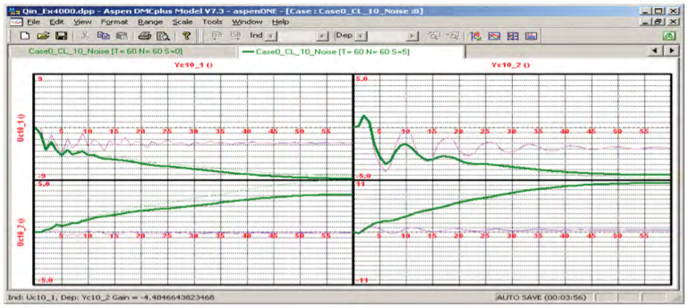

*Identification with PRBS-Test. FIR ID: Closed-Loop with 10% Noise.*

# Adaptive Process Control

In late 2011, researchers at Aspen Technology developed a new switching control algorithm that has the ability to control the process by making step-like changes to the MVs and stabilizing the process even if significant model mismatch is present, while also preventing the LP targets from flipping around. It is intended to work as a robust controller where minimal MV movement is made while active process constraints (active CVs) are kept at or close to limits. Disturbances are still rejected as needed, limit changes are followed, changes to feed forward signals are taken into account, and LP cost function changes are used immediately—all without the steady-state targets jumping around or the dynamic move plan engine making unacceptably large changes to the MVs.

In this way, the controller maintains process stability while also performing very small perturbation tests that generate closed-loop data. The data is of adequate quality which allows us to fit sufficiently accurate models via the new closed-loop Subspace identification algorithm. The process unit can be operated close to the point desired by the operations staff, even if the LP costs are not entirely correct. The LP can be setup to push a specific set of constraints via the Economic Relaxation parameter and MV cost ranks. Adaptive Process Control can be setup to ensure that product giveaway and feed rate reductions are minimized. In all scenarios, the engineer has the ability to define the degree of trade-off between optimal control and optimal test data. Moving from traditional sustained value tools to Adaptive Process Control requires more than robust controller behavior. To create an effective continuous process for updating models requires several pieces of innovation:

1. Automation to replicate human judgment when editing test data to remove bad values

2. The ability to detect PID loop status in conjunction with data editing

3. Economic Relaxation to provide room for the modest process perturbation while still maintaining optimizing control

4. Automation to generate candidate models

5. Methods to automatically determine the quality of candidate models

Adaptive Process Control combines the power of switching the controller with improved data pre-processing and identification technologies—enabling a new way of maintaining APC applications.

# Some Breathing Room

There have been many cases where controllers were not properly maintained and were eventually turned off by the operators. Peaks in the APC engineer's workload and a host of other reasons could result in extended periods of time going by between the onset of poor controller performance and corrective maintenance. If the condition is allowed to persist for too long, the operators will turn off the controller. Given that we do not always have the opportunity to immediately work on the controller, we need a way for the controller to change from an optimizing controller to a robust controller and thereby allow the application to continue to operate until the model updates can be completed.

While controlling the unit, we can also use this discrete robust control mode as a benign background perturbation mode to collect free test data while the unit is approaching optimal process operation. To do this, the MV sequence is made up of small GBN (Generalized Binary Noise) sequences plus discrete staircase-like MV movement when the Aspen DMCplus engine turns on to deal with CV excursions beyond the CV limit.

The main benefits of this approach include:

- Robustness: The discrete switching control approach is highly robust, especially with regard to co-linearity/RGA errors. In several simulations, RGA sign errors in weaker relationships could be tolerated. Of course, RGA sign errors in the dominant (strong) relationships should be avoided. Individual gain errors of 5-10x can often be tolerated. Model errors of this magnitude in the dominant relationships can be removed from the “seed” (starting) model with just a few days of test data.
- High-Quality Test Data: The data generated in this way contains only minimal amounts of feedback correlation and is therefore suitable for Subspace model identification, even if significant disturbances are present and even if the signal-to-noise ratio is well below 6:1. The latest version of the aspenONE APC Subspace identification algorithm is “closed-loop capable” and can handle up to 80% closed-loop data.
- Graceful Degradation Capability: The new controller algorithm underpinning Adaptive Process Control provides a way to immediately stabilize an Aspen DMCplus controller that is exhibiting model mismatch via a simple mode change in the web interface. This provides a graceful degradation capability in Aspen DMCplus, so it can handle slow or abrupt changes in process dynamics. The controller can then run in Adaptive mode for several weeks or even months, while the software sifts through the data to remove undesirable segments and produces candidate models without the need for human intervention. When the model matrix converges, the RGA repair work can be done via SmartAudit and the Aspen DMCplus controller can then be re-commissioned, getting us back into profit hungry mode.
- Moderate-to-Good Process Optimization Capabilities: In Adaptive mode, the process unit can be driven fairly close to the desired operating point even if the LP costs and gains are not entirely correct. The LP can be setup to push towards a specific set of constraints via the Economic Relaxation parameter and MV/CV cost ranks. In this way, Adaptive Process Control can be configured to ensure that product giveaway and feed rate reductions will be negligible, so that the unit will be seen as operating better than before.
- More Productive and Less Time Intensive: Full time engineering supervision will not be required for more than the first few days when the controller is in Adaptive mode. Once the strong models have been updated with new model curves via the web interface, the MV and CV limits can be opened up. Once a week, the engineer can have a look at the new model and decide if an update is needed.

 

# Conclusions

The synergy between innovations in model identification, automated testing and controller operating modes are creating opportunities to bring new efficiencies to the process of sustaining APC controllers. These innovations address some of the long-standing challenges facing APC practitioners: dealing with closed-loop data, managing APC solutions with model mismatch until repairs can be completed, and reducing the operational impact of step testing.

With these advances, we now have the technology to begin constructing continuous improvement processes in lieu of the traditional sustained value workflows. We have the ability to switch between optimizing control and constraint control. We can gently perturb the process while in constraint control to gather usable process data, and we have improved model identification software to produce good models using closed-loop data and the data from the background perturbation of the process. By unifying these new capabilities with integrated interfaces and workflows within the software, we're delivering Adaptive Process Control – an iterative, continuous process for sustaining the value of APC applications.

# Appendix I

## How Does Adaptive Process Control Work?

When in Adaptive mode, the optimization algorithm recalculates the steady-state (LP) objective function at every time step and stores this cost function ($J_{optimal}$). It then calculates the cost function at the current MV targets ($J_{current}$) with minimum move setting (constraint control), compares the two cost function values, and calculates $\Delta J = J_{optimal} - J_{current}$. If $\Delta J$ is larger than the user-specified tolerance, the Aspen DMCplus engine will generate a move plan and push the process towards the new optimal solution. However, most of the time it will not implement the new MV targets since the two solutions are usually within the specified tolerance.

In this diagram, $y_1$ and $y_2$ are the two active CV constraints, and $u_1$ and $u_2$ represent two particular MVs.

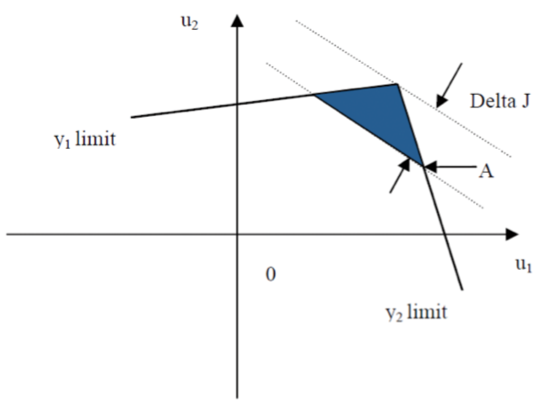


The technology ensures that the two CVs will be kept inside of the blue triangle. In effect, an additional cost constraint is imposed so that the cost function will not degrade more than the user-specified $\Delta J$ (Delta J). When inside the blue triangle, the engine will generate step moves for the requested manipulated variables, similar to the multi-test in the traditional SmartStep with the exception that the step move will not violate the user-specified $\Delta J$ constraint.

## How Does the Algorithm Perform in Practice?

Below is a simulation using Aspen DMCplus with zero model mismatch and where all the CVs have wide ranges, so we rely on the LP to determine which MVs should move up or down. (We will introduce very large model mismatch at a later point in the discussion.)

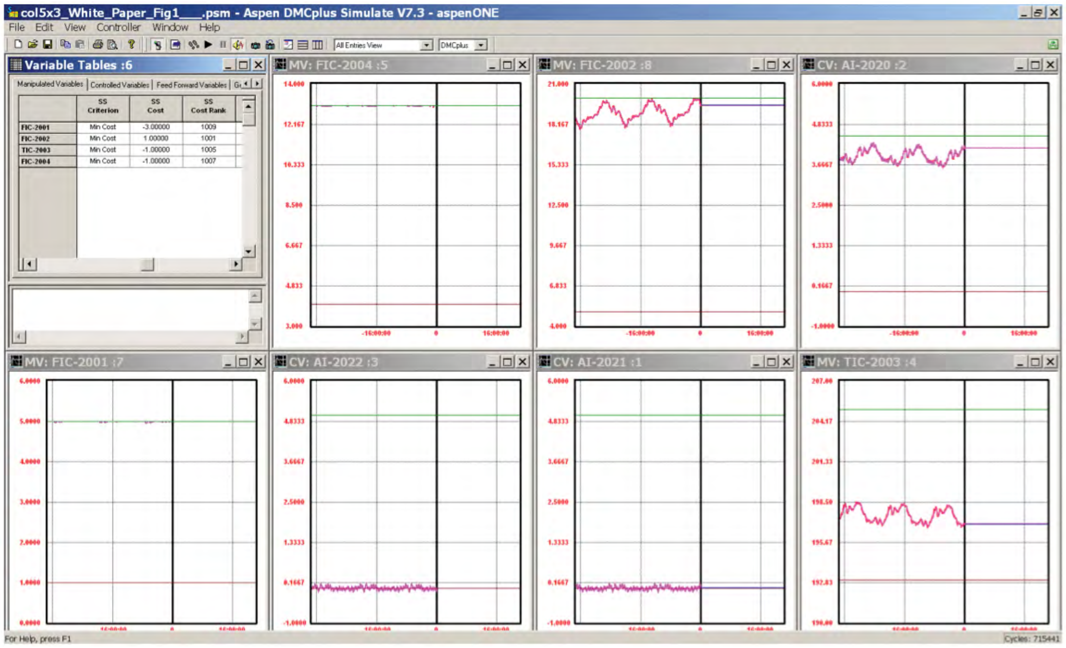

Note that the low limit is active for CV:AI-2021 and CV:AI-2022, while MV:FIC-2001 and MV:FIC-2004 are pushing towards their high limits (see below).

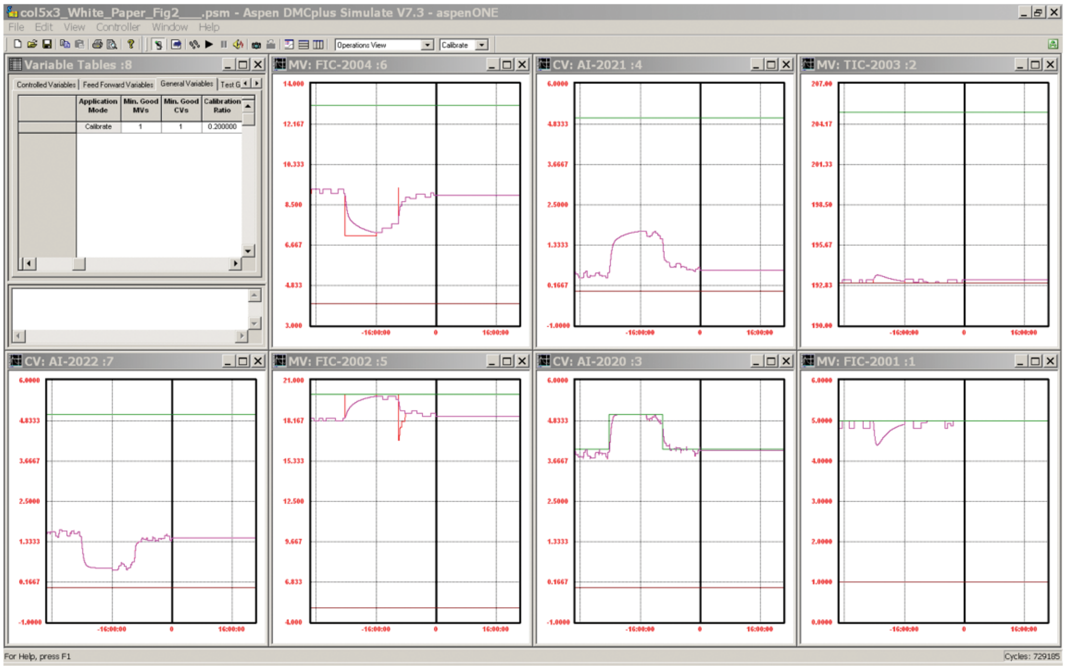

Also note how MV:FIC-2002 goes through the same repetitive sequence as it rejects this simulated disturbance, which repeats every 16 hours. The disturbance is made up of a ramp that changes direction every 8 hours, plus random noise.

Now let's switch over to Adaptive mode and test for the following:

- Will the controller still push the unit towards the same active constraint set?
- Can it respond appropriately to CV limit changes?
- Will it correctly deal with unexpected disturbances?

To more clearly see what is going on, we will start with a noise-free simulation. A second example is later presented showing how the new engine responds to noise and large disturbances. We set the Adaptive mode Calibration Ratio to 0.20 (or roughly 20%, a fairly large value), used the same tuning and limits, maintained zero model mismatch, and switched the application from an Aspen DMCplus operation into Adaptive mode - see the figure below.

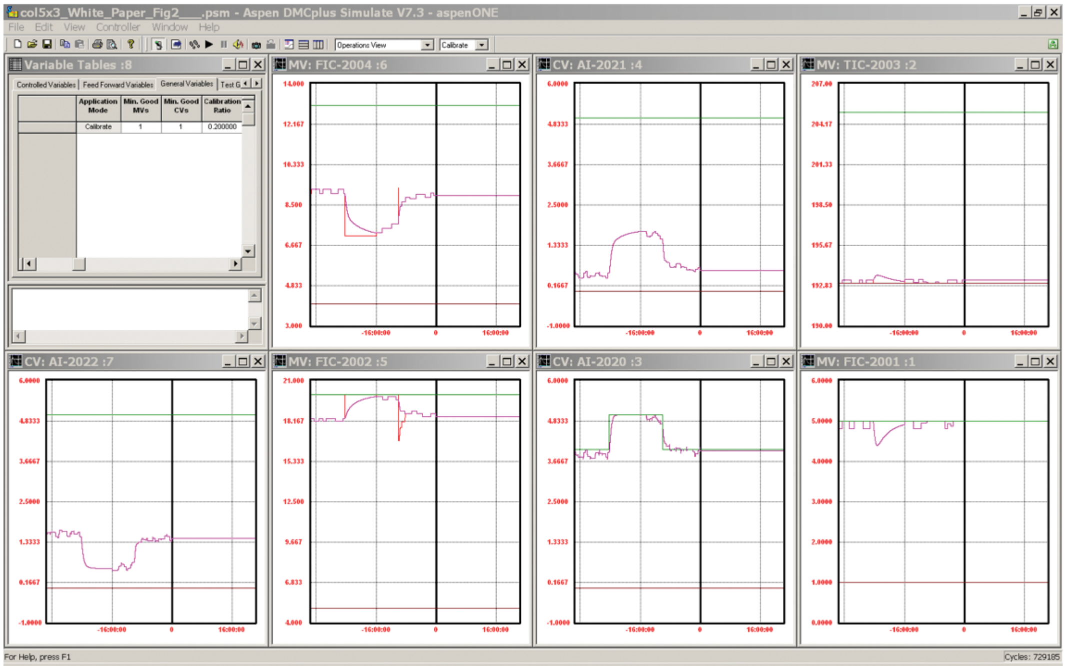


Note that the operator high limit for CV:AI:2020 was increased from 4.0 to 5.0 (green line) at the start of the simulation. The step-like movement in the red MV LP targets indicate that the Aspen DMCplus algorithm turned on and a set of Aspen DMCplus control actions took place to drive the CV to the new high limit.

 

A few hours later, the algorithm changed from an Aspen DMCplus operation back into an economically constrained multi-test operation, with small MV steps being introduced again. The CV high limit was then changed again from 5.0 to 4.0. The red MV LP targets indicate that the Aspen DMCplus engine turned on for a short period of time, and once the cost function was close to optimal, the multi-test algorithm took over to make a series of step-like moves that drove the CV towards the new CV limit. The algorithm could now make a larger-than-usual step in FIC-2004, since there is now more room available than before.

If we used a much smaller calibration ratio, or if we made a larger CV limit change, ΔJ would be big enough and the engine would switch over to the Aspen DMCplus control mode for 0.5 * time to steady-state (TTSS).

Note that the active CV:AI:2020 is moved just inside its CV limits, and just barely touches the CV high limit during the largest excursions (due to a pronounced inverse model response). This is typical behavior while in Adaptive mode. If you believe that these small changes in average MV and CV values will lead to appreciable giveaway losses, then simply open up the MV and CV limits by a small amount. In practice, the steps are so small that they will be barely visible above the normal process noise and disturbance effects.


As mentioned before, the average value of the CV is shifted down in order to be inside the CV low and high limits, compared to Aspen DMCplus which would have an average CV limit equal to the CV limit (about 50% of the values would be above the limit, and 50% below the limit). It aims to be slightly conservative. Notice how Adaptive mode causes the controller to actively follow the CV limit changes, always trying to stay on the safe side of the active CV limit. The MV movement appears to be very similar to what you would expect from a controller with limited MV resolution. In fact, in this configuration, multi-test is setup to act like a very benign discrete amplitude switching controller, with fairly small MV steps imposed on top of an MV sequence that follows a staircase trajectory.

 

What are the benefits of using a discrete controller?  First, the MIMO aspects of the model (RGA numbers and gain ratio differences) can be substantially inaccurate if the process unit has degraded since the original step test was done. The same is true early on in a new Aspen DMCplus project when we have limited test data. The algorithm was designed to maintain stable control of the process unit in the presence of very large model errors. Simulations were constructed that contained gain errors as large as 10x – large enough that the Aspen DMCplus controller went unstable due to LP flipping, yet Adaptive mode maintained stable operation of the unit.

Even if the individual strong model curves (which are the only ones that matter while in Adaptive mode) have significant error in terms of shape and gain value, the algorithm will be able to handle it. All we assume is that the sign of the strong models (on typical move scaling) are correct, and within one order of magnitude of the true value. Weak models may be absent or may even have the wrong sign. Adaptive APC is clever enough to largely ignore weak relationships, expecting them to be highly inaccurate. The algorithm avoids having to invert 2x2 or larger square matrices, and we have found that the co-linearity repair is not yet needed and the seed model may be very sparse. Once we are ready to go back into the Aspen DMCplus Control mode, the co-linearity analysis and repair tool should be used to fix the RGAs down to a threshold of at least 10, preferably as low as 5.0.

## How Does Adaptive Process Control Handle Process Disturbances?

So far, we have only looked at a noise-free simulation. If we add major disturbances to the simulation (noise plus prolonged periods of strong drift), we get the following result below while in Adaptive mode.

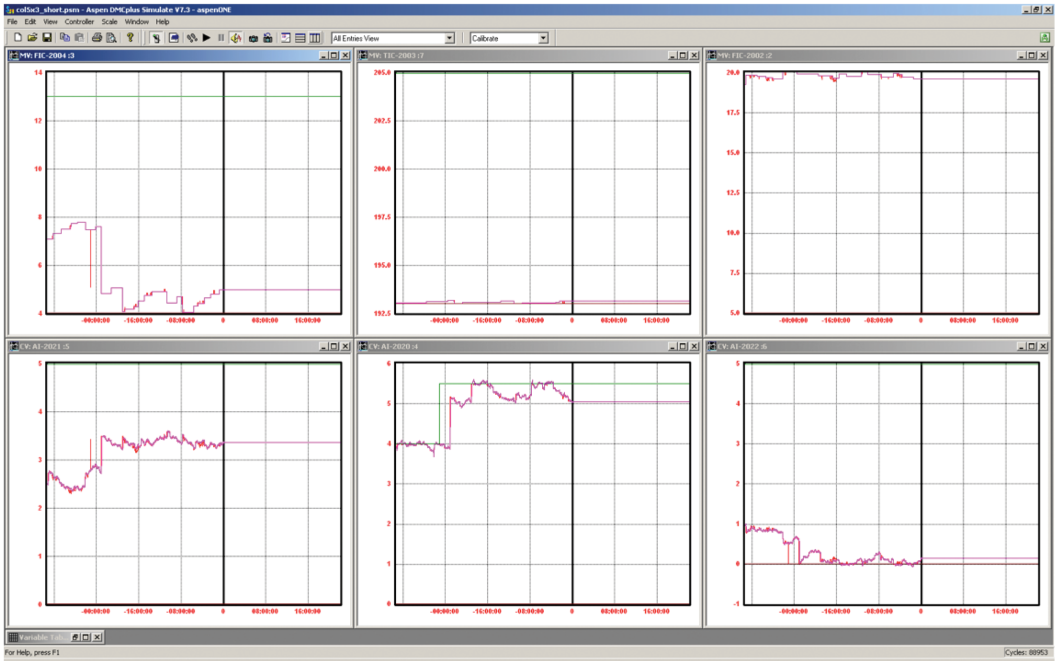

As before, note the small blip in the MV LP target (in red) at the point in time when the CV limit change was done. The value of the LP targets at this point indicates where the Aspen DMCplus controller would have moved the MVs and CVs, had it decided to turn on. Once again, the ΔJ was too small to change the mode.

 

The limit changes are followed after a two hour delay, reaching the LP targets calculated at the time of the limit change. Once at the active constraints, the CVs tend to touch the CV limit only during worst case disturbances. In effect, the algorithm is trying to iteratively determine MV targets that keep the CVs just inside the CV constraints.

# Appendix II

## How Does Adaptive Process Control Deal with Extreme Model Mismatch?

In a recent forum post on LinkedIn.com, the participant asked an interesting question, "You only want to revamp a controller with a poor model. A poor controller would be flipping the LP targets frequently, the MVs could be cycling, or the LP could be pushing in the wrong direction. How can you use a poor model to safely step test a process unit?"

The real question being asked here is: Does bootstrapping work? Can we start with a low accuracy model, with many missing model curves, perhaps even some sign errors, and get Adaptive APC going reliably so that it can generate new data that we can use to re-identify the model matrix?

Let's start with very substantial model mismatch (in fact, excessively large model mismatch), and see what happens to Aspen DMCplus. The true process response is shown in blue below, while the controller model is shown in red.

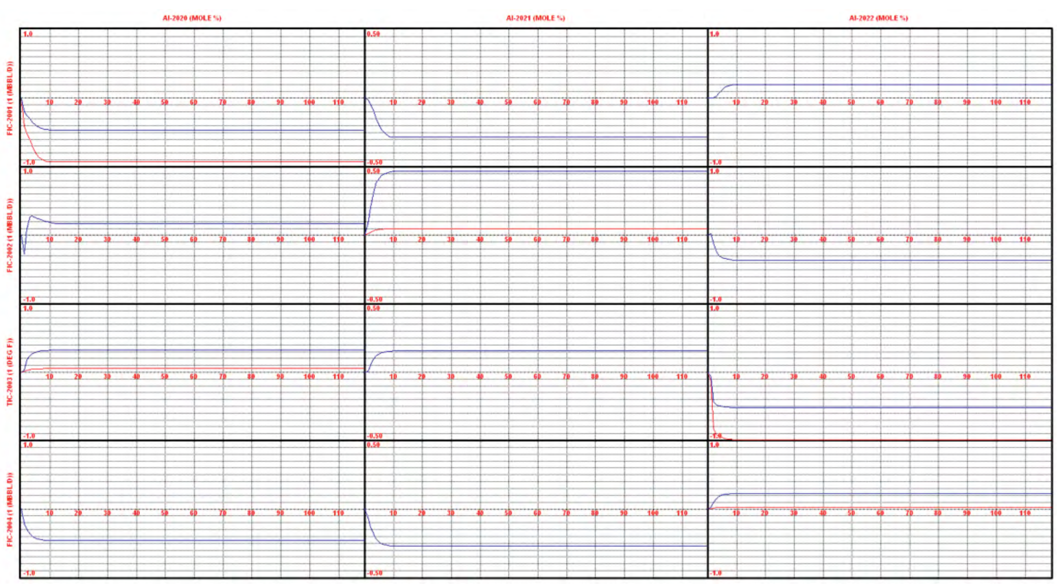


Note that this is a distillation column and there is a model curve in every cell (full density). The controller model matrix is shown in red and has very few model curves (low density). With the exception of the tray temperature MV, only one model curve is present per CV, and very large model errors are present (between 2x and 10x).

The controller model has been setup with very low model density, as shown in the figure below.

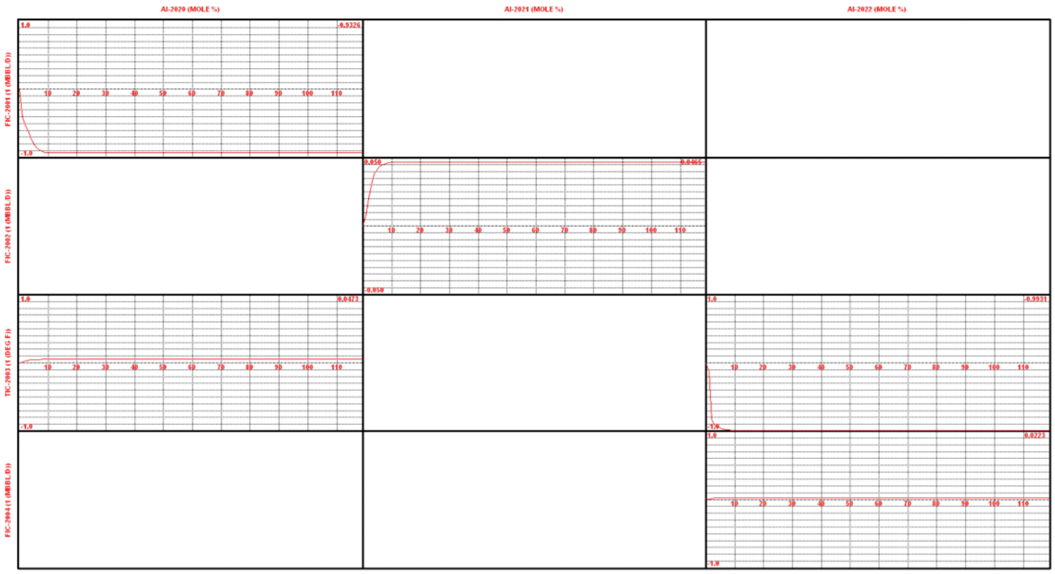

All the feed forward models have been removed. These model errors were generated via random gain multipliers.

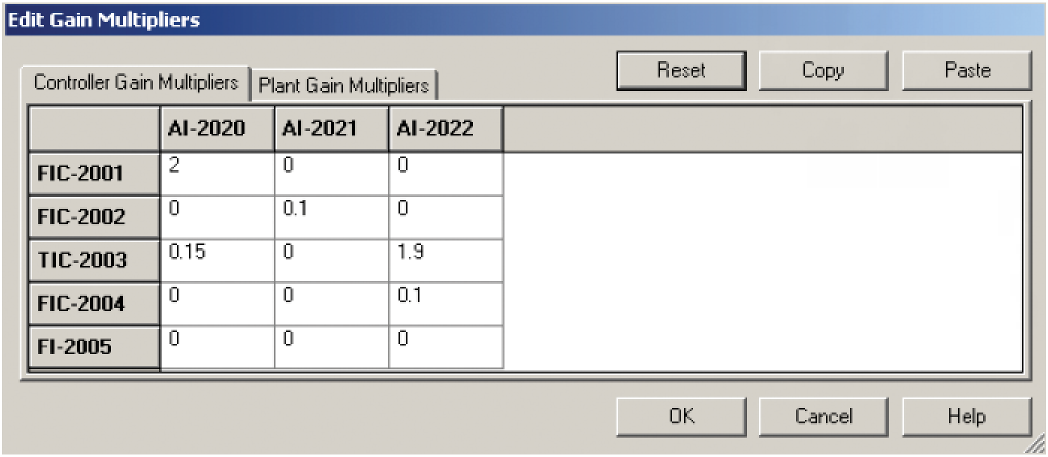



| | AI-2020 | AI-2021 | AI-2022 |
|---|---|---|---|
| FIC-2001 | 2 | 0 | 0 |
| FIC-2002 | 0 | 0.1 | 0 |
| TIC-2003 | 0.15 | 0 | 1.9 |
| FIC-2004 | 0 | 0 | 0.1 |
| FI-2005 | 0 | 0 | 0 |

As can be seen above, the gain errors vary between 0.1 (10x too low, i.e. worst case direction) and 2.0 (2x too high). The majority of the model curves have been removed (zero GMULT values).

At 16:00 we introduced this very substantial model mismatch. Let's see what happens during the Aspen DMCplus operation below.

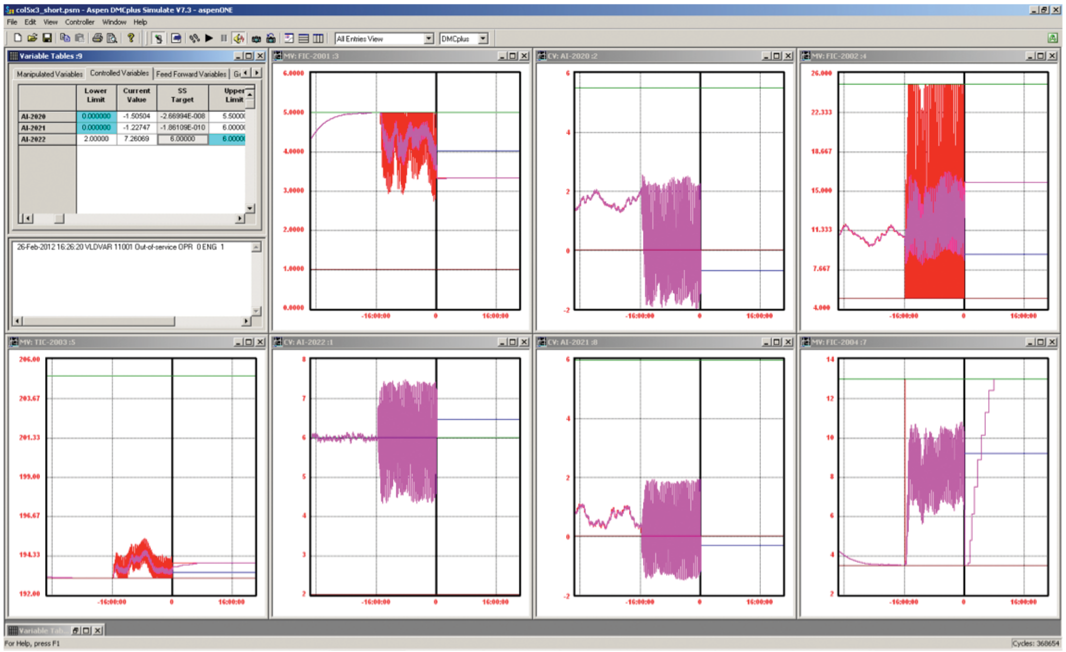


Note how the Aspen DMCplus controller goes unstable, which isn't surprising given the excessive level of model mismatch. The MV cycle grows exponentially until the MV LP targets hit the MV limits, then it settles into a constant amplitude cycle. The main source of instability is MV:FIC-2004, which has an LP target that flips between low and high limits every few minutes.

The next question is whether Adaptive Process Control can use a model this poor and remain stable and well-behaved, with small CV limit excursions. The same model mismatch is used in the subsequent simulation, and the mode is changed from "Adaptive" mode to "Aspen DMCplus" mode.

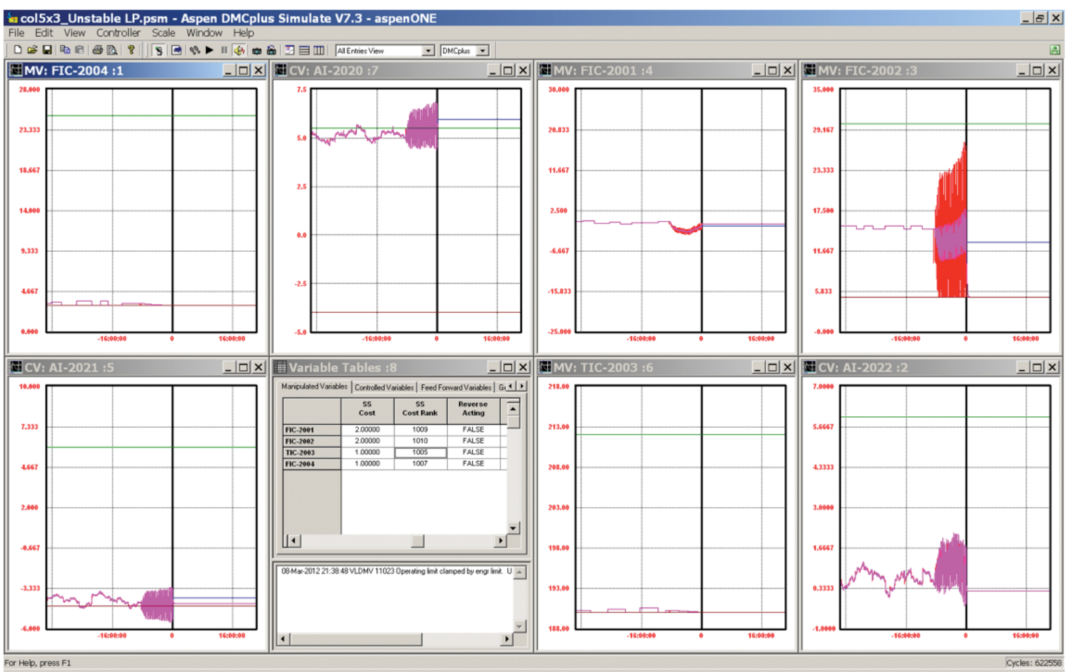

It is clear that the first 24 hours of the simulation shows Adaptive Process Control in action, the MV movements are benign and are not growing, and the active CV:AI:2021 just touched the CV low limit a few times with little overshoot. Once the application is switched back into Aspen DMCplus mode, the controller goes unstable and the LP target starts cycling badly. This simulation clearly demonstrates that Adaptive APC can handle extreme levels of model mismatch, far more than what a standard LP driven MPC formulation can tolerate.

# Appendix III

## How to Set the Adaptive Calibration Ratio

If we set the Adaptive Calibration Ratio to a small value, the switching controller will stay very close to the optimal constraint set, but the MV steps will get highly correlated with the unknown disturbance. This effect is called feedback correlation and is the main reason why the model ID fails in the presence of pure closed-loop data (with no external perturbations of any kind).

To make it hard, let's use the large ramp/triangular disturbance we used before and tell the controller that the process gain is only 15% of what it really is (model mismatch of 6.7x, which is a very large gain error).

Let's consider a small ratio like 0.01 and see if we can identify the models.

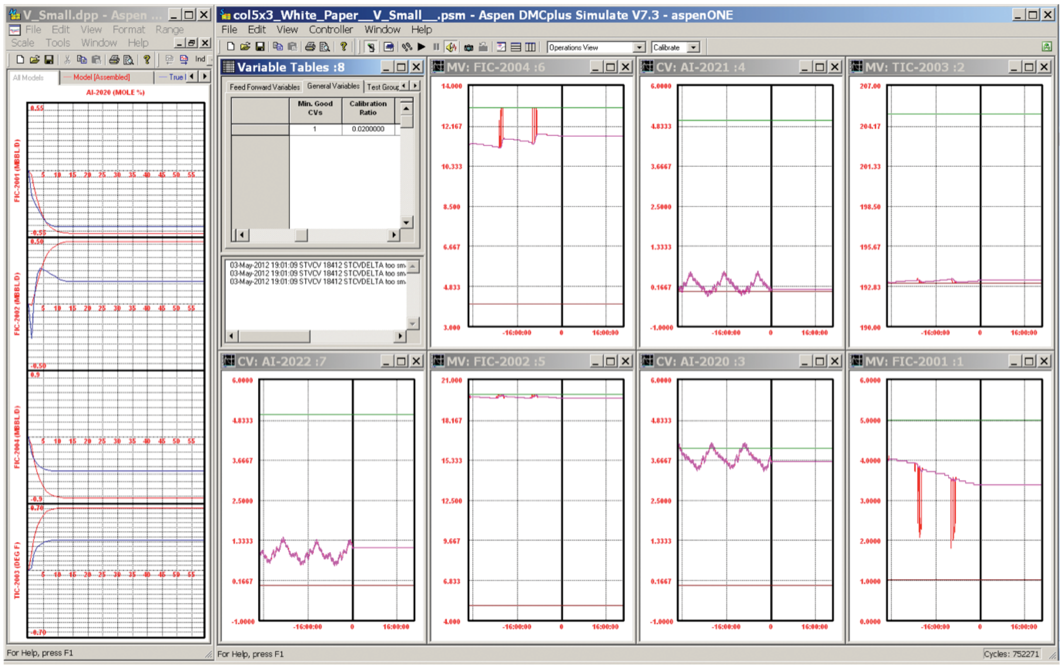


 

Note how the controller has to keep ramping the MV to prevent the CV from dropping below the CV low limit. This half of the cycle almost looks like pure closed-loop Aspen DMCplus control data (in fact it is slightly better than that, but not much). The second half of the cycle is made up of a series of larger step-like staircase moves. The CV stays very close to its LP target, but the MV takes on almost the exact shape of the triangular disturbance and is therefore strongly feedback correlated. The model ID results are shown on the left. Rather surprisingly, the new closed-loop version of the Subspace identification algorithm has managed to identify model curves that are within 20% of the correct gains. This is very impressive for such a short data set (and such a difficult disturbance).

Of course, if the disturbances are random and we collect multiple weeks of data, we will eventually identify a much more accurate model than shown here. The theory predicts that if the process is linear and if we reduce step size by 10x, we will get about the same model accuracy if we collect 10x more data. With a Calibration Ratio this low (0.01), we cannot be sure we will get models of adequate accuracy, even if we collect months of data. Values this low is not recommended.

Note that the MV steps still need to be bigger than the stiction and hysteresis effects, otherwise these nonlinearities will dominate the data set and the model ID results will not be good.

The linearity is only approximately maintained if the step sizes aren't larger so that the PID outputs will saturate every time, but still big enough so that the valve stiction and hysteresis effects are not dominant in the data set.

Of course, ratios around 0.05 to 0.15 would provide for the opportunity to make steps in opposite directions and increase signal-to-noise ratio. The model ID results improve even further and less data is then needed to converge the model matrix. If we use a ratio of 0.1, we get bigger steps and even better model ID results (see below).

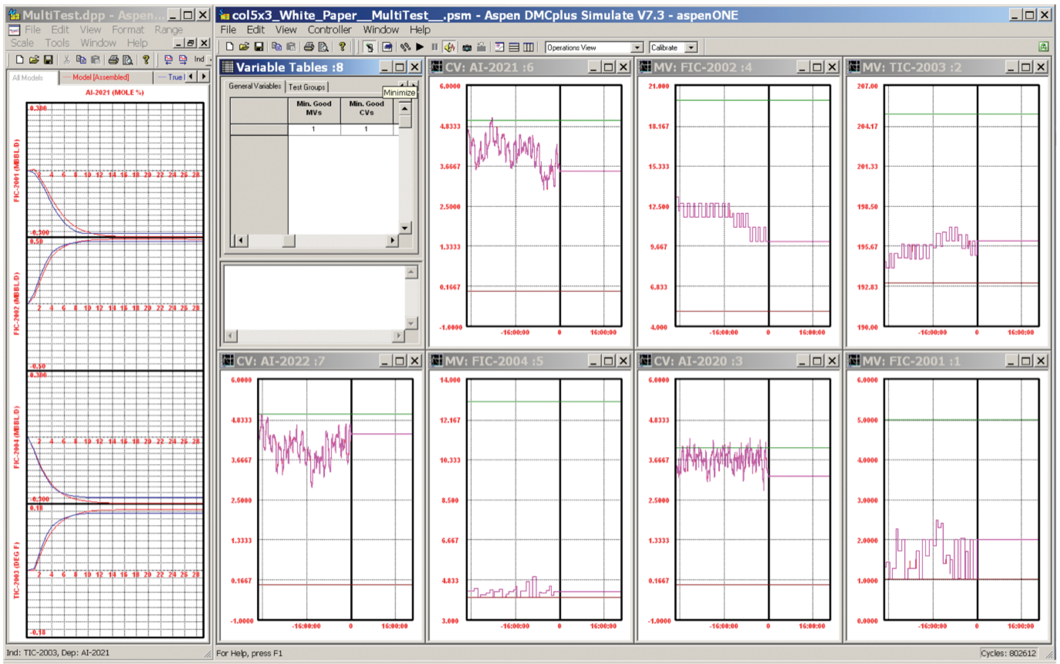

Note that the MV ramps still oppose the ramp-like disturbance, so there is still some feedback correlation in the data. However, compared to the previous example, the MVs do show direction changes, so some small steps are being made, breaking the correlation. This greatly reduces the feedback correlation and significantly improves the model ID results, as shown. We typically need one and a half to three weeks of data if we select a ratio of 0.1. This is the maximum allowed, which would result in the following shown below.

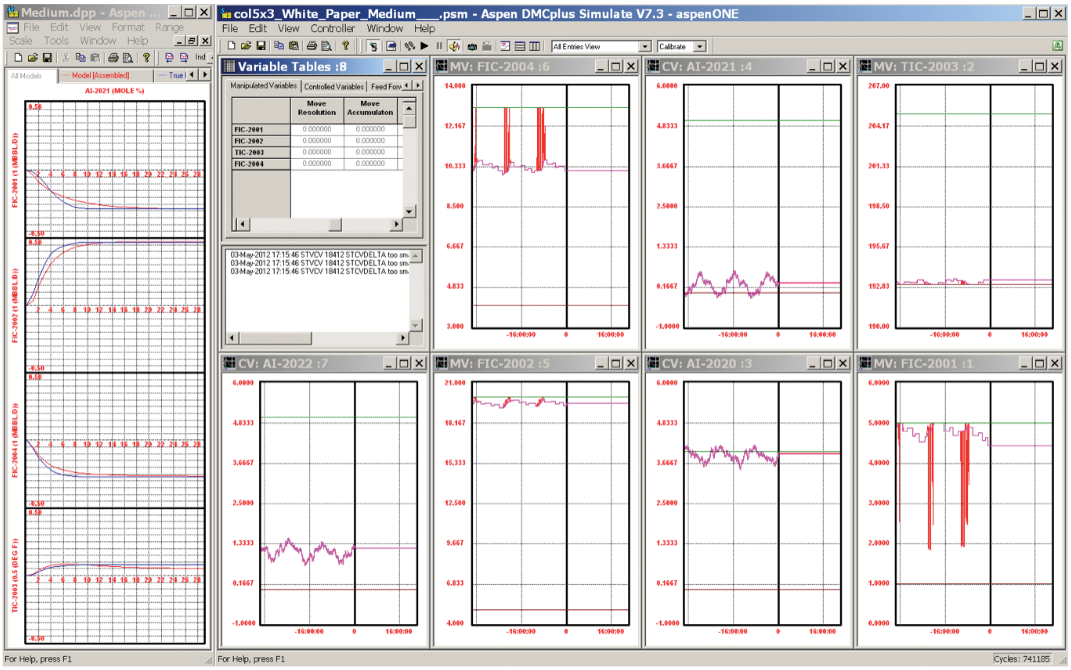


The Adaptive Process Control algorithm is still pushing the active constraints, but the signal-to-noise ratio has improved significantly. The model ID results are almost perfect, even though this is less than two days of test data.

At a ratio of 1.0, we can expect the strong models to converge in two to three days. The moderately strong models (also required) will take one to two weeks. The weak models can be ignored.

AspenTech is a leading supplier of software that optimizes process manufacturing — for energy, chemicals, engineering and construction, and other industries that manufacture and produce products from a chemical process. With integrated aspenONE® solutions, process manufacturers can implement best practices for optimizing their engineering, manufacturing, and supply chain operations. As a result, AspenTech customers are better able to increase capacity, improve margins, reduce costs, and become more energy efficient. To see how the world's leading process manufacturers rely on AspenTech to achieve their operational excellence goals, visit **www.aspentech.com**.


**Worldwide Headquarters**

Aspen Technology, Inc.
20 Crosby Drive | Bedford, MA 01730 | United States
phone: +1-781-221-6400 | fax: +1-781-221-6410 | **info@aspentech.com**

**Regional Headquarters**

Houston, TX | United States
phone: +1-281-584-1000

São Paulo | Brazil
phone: +55-11-3443-6261

Reading | United Kingdom
phone: +44-(0)-1189-226400

Singapore | Republic of Singapore
phone: +65-6395-3900

Manama | Bahrain
phone: +973-13606-400


For a complete list of offices, please visit **www.aspentech.com/locations**