

Simplify APC Tuning and Configuration

Advanced Process Control

Tushar Singh, APC Product Marketing, Aspen Technology, Inc.

Lucas Reis, APC Product Management, Aspen Technology, Inc.



Introduction

Typically, one of the most technically challenging parts of deploying and maintaining advanced process control (APC) applications is the tuning of the steady-state optimizer. The APC engineer needs to understand the economic objectives of the process unit and then determine the appropriate optimizer tuning (LP costs) to drive the controller to achieve these objectives. For many process units, the knowledge of which constraints need to be pushed and the manipulated variables' (MVs) preferred direction and optimization sequence to maximize profits is already known to operations and APC engineers (for example maximize feed first, then maximize products and then try to minimize energy consumption). The APC engineer uses his experience, skillset and in-house tools to devise a set of LP tuning that forces the optimizer to achieve the desired controller behavior. Often, the process of determining LP costs involves several iterations and the use of simulations to make sure the controllers economic objectives are achieved under several different scenarios (constraint sets). With traditional APC steady-state optimizers, the problem becomes increasingly challenging as the number of variables or process interactions increase. Another issue is if models are updated, the LP tuning needs to be revisited to verify that the controller is still achieving the intended objectives.

The complex nature of tuning the steady-state optimizer in a traditional APC controller demands some degree of experience and efforts to setup and modify. To reduce this complexity, a whole new approach to APC steady-state optimizers is required. Instead of using the traditional "one objective function for the entire controller" approach, using a Sequential Multi-Objective Optimization (SMO) technology — patent pending — enables the user to directly specify the operational objectives thereby drastically reducing the efforts required to set up the steady-state optimizer. This paper describes the use of this novel technology to help APC engineers significantly reduce the efforts required to tune APC controllers.

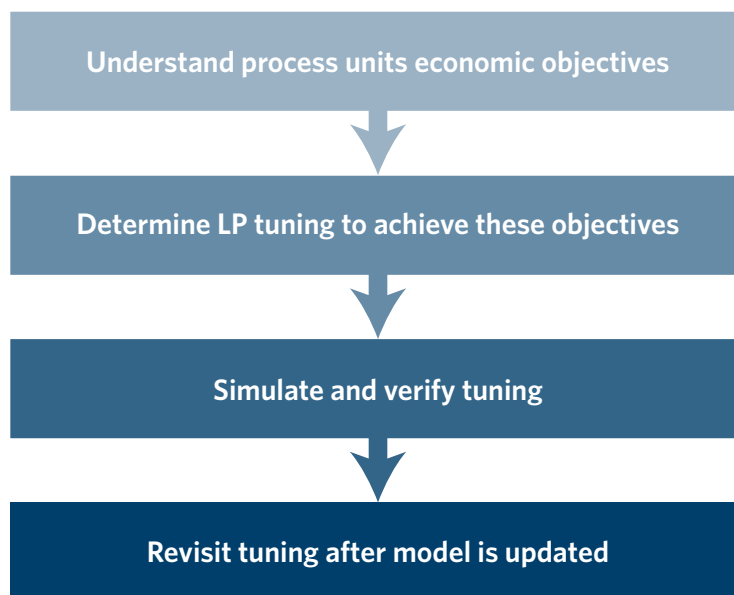


Figure 1: Traditional Path to APC Controller Tuning

Challenges with Steady-State Optimizer Tuning

A correctly configured steady-state optimizer is key to delivering the benefits associated with an APC project. The steady-state optimizer calculates targets for manipulated variables (MVs) and controller variables (CVs) and is responsible for the “where do you want to go?” part of the control problem (Figure 2). These targets are then passed to the dynamic optimizer that calculates the move plan for the MVs. The dynamic optimizer is responsible for the “how will you get there?” part of the control formulation. The steady-state optimizer is used to find the most economically favorable solution within the operating constraints of the MVs and CVs. It uses LP/MV costs to drive the process unit towards maximum profitability, while remaining within the operating constraints. Ideally, these MV costs should represent the real unit costs (e.g. \$/kg). In reality, we rarely know costs of all MVs in the controller, for example we may know the cost of our feed, but the cost of intermediate streams or process variables may be unclear. Accurately knowing the costs is not the only challenge associated with setting up the optimizer. Even if we can accurately calculate the cost of all the variables, the use of these costs necessitates the use of very accurate APC models.

Clearly, using real MV costs in many cases is not practical. Because of these challenges, most APC engineers do not use real MV costs but instead treat them as “LP tuning” factors.

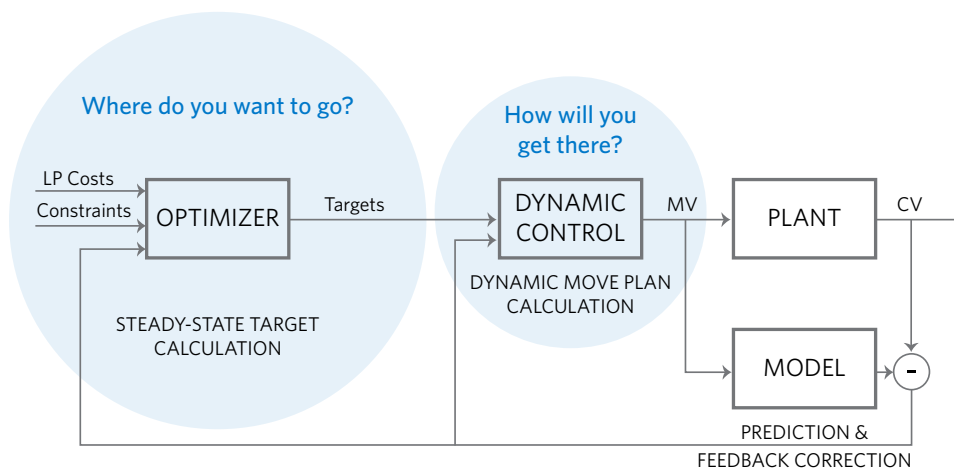


Figure 2: APC Engine Setup

LP or MV costs are used by weighing in the objective function to trade-off MV targets. MVs with “relatively” higher costs are driven lower. An emphasis on “relative” is important. Since traditional APC applications use one objective function, the absolute value of MV costs have little meaning. What’s important is the relative MV cost factors. The final solution is affected by a combination of these relative MV costs, MV/CV constraints and model gains. This means that as the model matrix becomes dense, it becomes increasingly challenging to come up with a set of LP cost factors that achieve the intended objective. Most APC engineers use iterative methods and/or external tools to initially determine these LP cost factors. This is a time-consuming process and requires experience and efforts to set up. Since the steady-state solution is also a function of the model gains, it becomes important to revisit the LP tuning after a model is updated to ensure the intended controller objectives are not affected.

Below is a summary of the challenges APC engineers face in the setup of traditional optimizers:

- Cost factors are not readily available and demand accurate models, hence they are set by trial and error to achieve the desired controller objectives.
- Operators and engineers have a preference for which constraints should be pushed during normal operations. The cost factors are then tweaked until the optimizer results are consistent with this strategy.
- Whenever controller objectives or models change, the costs factors need to be revisited to ensure the desired controller behavior.
- For a less experienced APC engineer, the controller configuration can be a challenging task.
- LP tuning requires significant efforts, especially when the model becomes larger and/or denser.

A Novel Approach to the Steady-State Optimizer

Most of the challenges discussed previously can be addressed if traditional APC optimizers were modified to have not one, but multiple steady state-objective functions for an APC controller. A Sequential Multi-Objective Optimization technology (SMO) can be used to explicitly specify economic objectives directly into the controller rather than costs and thus eliminate the need to tune LP costs. We will demonstrate the ease of setting up the optimizer using the new multiple objective optimizer approach through an example that follows. See how the new technology called Smart Tune, available in Aspen DMC3, is different than previous generations of APC optimizers.

Example

APC has been successfully and widely deployed on many different processes in the process industry. One such process is the crude distillation unit (CDU) in a refinery. Although we will demonstrate the Smart Tune technology using a simple crude distillation unit (CDU), it is in no way restricted to any process. Smart Tune can be used anywhere APC technology can be applied.

Let's assume that we have a simple CDU unit represented by the following process flow diagram in Figure 3.

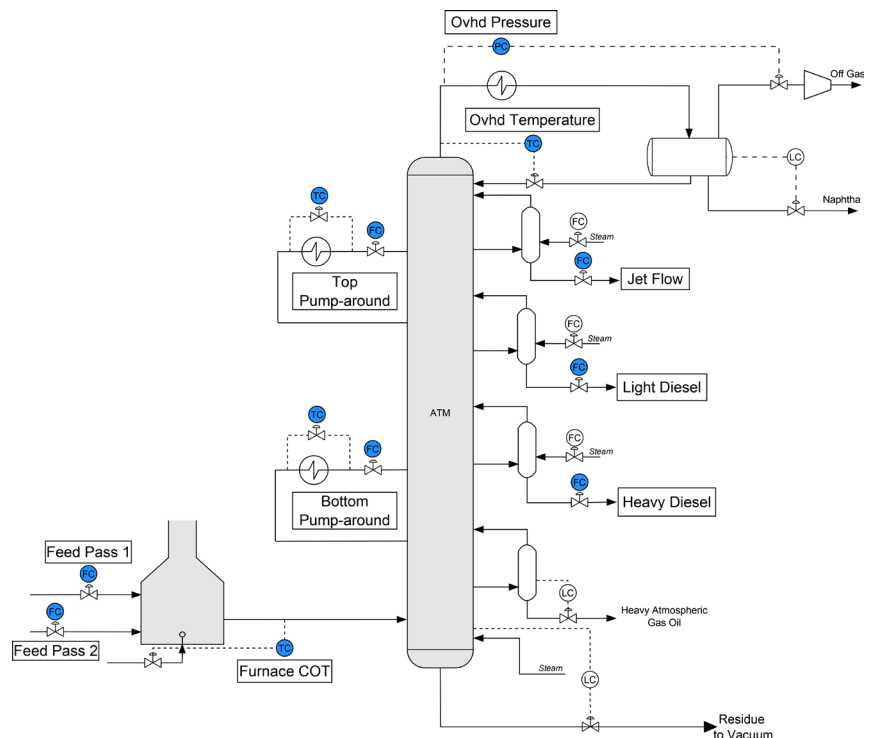


Figure 3: Crude Distillation Unit APC

Defining Controller Objectives

The first step in building an APC controller is to define the objectives of the controller. Let's assume that the CDU APC controller example above has the following objectives for its MVs.

MVs with clear economic optimization direction:

- Feed Pass Flows 1 & 2 (Maximize)
- Furnace COT (Maximize)
- Jet Flow (Maximize)
- Light Diesel (Maximize)
- Heavy Diesel (Maximize)
- Overhead Pressure
(Do not move for optimization/Minimize movement)
- Overhead Temperature (Minimize)

MVs with no preferred optimization direction:

- Top Circulation Flow
- Top Circulation Temp
- Bottom Circulation Flow
- Bottom Circulation Temp

Priority	MVs
1	Feed Pass 1 (Maximize)
1	Feed Pass 2 (Maximize)
2	Furnace COT (Maximize)
3	Jet Flow (Maximize)
4	Light Diesel (Maximize)
4	Heavy Diesel (Maximize)
5	Overhead Pressure (Min Move)
6	Overhead Temp (Minimize)
7	Top Circ. Flow (No Preference)
7	Top Circ. Temp (No Preference)
7	Bottom Circ. Flow (No Preference)
7	Bottom Circ. Temp (No Preference)

Figure 4: MVs Priority and Preferred Direction

Although we have defined the MVs preferred optimization direction, there is still a question of defining which MVs achieve their preferred direction first. For example, we know we should maximize Jet Flow and Light Diesel. We also know that increasing Jet Flow will result in some decrease in Light Diesel. Therefore, the question remains whether we should maximize Jet Flow over Light Diesel. For the sake of this example, let's assume that economics or scheduling constraints dictate that we maximize Jet Flow over Light Diesel. Note that we still want to maximize Light Diesel, but want to achieve this by not sacrificing any Jet Flow. For the sake of demonstrating the technology, let's also assume that we would like to trade-off Light Diesel and Heavy Diesel using their true costs. After going through this exercise for all the MVs, we have the priorities and preferred direction defined for our MVs, shown in Figure 4.

Notice that there are some MVs that have the same priority. Feed Pass 1 and 2 should be maximized equally, i.e, there is no advantage of maximizing one over the other. Light and Heavy Diesel are in one group where we intend to use real costs associated with these MVs. Also notice that the MVs with no preferred optimization direction have the least priority.

Traditional Steady-State Optimizer Configuration

Now that we have our controller objectives defined, it's time to set up the steady-state optimizer in the controller. With a traditional controller we would need to come up with a set of LP costs that achieve all the intended controller objectives. The optimizer uses LP costs, model gains and constraints to trade-off all the MVs with each other in the same objective function. For this reason, tuning these LP costs is not very intuitive, especially if the model is dense and large. There is a lot of "interaction" between the LP costs since the MVs will compete not just based on cost, but also on shared CV constraints and model gains. For example, simply having a lower cost on Jet Flow does not guarantee that it would be maximized over Light or Heavy Diesel. This also means that it becomes very challenging to incorporate real costs of Light and Heavy Diesel. How would we make sure that the feed does not reduce for a CV constraint that Light Diesel could have relieved? In short, traditional controllers have no way of explicitly configuring individual MV priorities because they have only one objective function. The costs are the only handles available to achieve the desired controller objective. Setting up these costs is a time-consuming iterative process and needs to be revisited when models or objectives change.

How is the New Smart Tune Technology Different?

The new Smart Tune algorithm is different in the sense that its optimizer has multiple objective functions that are solved sequentially. This means that we can explicitly specify the controller objectives, i.e., priorities and preferred optimization direction of MVs directly within Smart Tune without worrying about LP cost tuning. This concept will be clearer once we compare the steady-state optimizer's objective function of Aspen DMCplus against Aspen DMC3. Figure 5 depicts the two objective functions.

In Aspen DMCplus, all the MVs are in one objective function. For this reason, as mentioned before, all MVs are traded off against each other using LP costs, model gains and shared constraints.

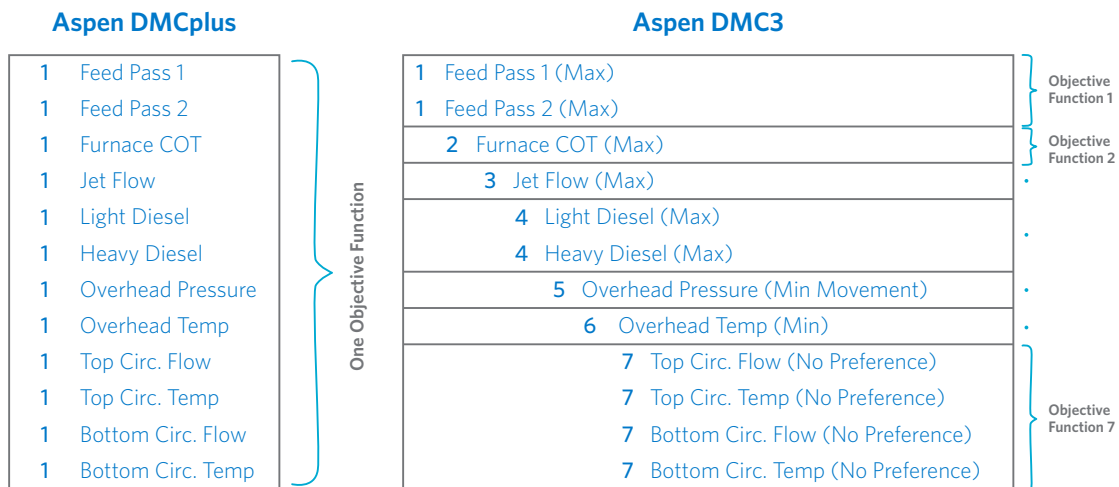


Figure 5: Objective Function Comparison

On the right in Figure 5 is the new Smart Tune objective function hierarchy in Aspen DMC3. As you can see the optimizer in Smart Tune technology is divided into multiple objective functions based on the priority groups that were defined in the controller objectives previously. All MVs in one priority group have their own dedicated economic objective function. The new algorithm solves the optimization problem sequentially beginning with the highest priority group. The highest priority group is solved by assuming that the remaining MVs in a lower priority group have no cost (free to move in either direction). After obtaining a solution for a higher priority, the algorithm optimizes MVs in the next priority group. It does this by forcing MVs in higher priority groups to not move when optimizing any MVs in a lower priority group. In essence, steady-state values of MVs in higher priority groups are fixed, while the solver optimizes a lower priority group. The algorithm continues until all priority groups are sequentially solved.

Here are some specific examples to further illustrate this new technology. In the CDU case above, let's say the solver is at a point where it is going to optimize priority group 3. Since Jet Flow is the only MV in that group, the optimizer will drive Jet flow to maximize (its preferred direction) assuming all other MVs with lower priority, i.e., Light and Heavy Diesel, overhead pressure etc., are free to move in either direction. The optimizer disregards preferred direction of all the MVs in the lower priority group. Also note that at this point, Feed Pass 1 & 2 and Furnace COT are fixed and are already maximized (their preferred direction). After the solver finds a Jet flow steady-state value, it moves on to optimize group 4. This group has Light and Heavy Diesel, therefore when the optimizer solves group 4 it uses the MV costs to trade-off Heavy and Light Diesel. Since the problem is broken down into a manageable set of trade-offs, we can now use real costs on these two MVs which are used to trade-off only these two MVs. The MV costs of Light and Heavy Diesel do not affect any MVs that are in a higher priority group (these are fixed during the optimization run). After group 4 is solved, the optimizer moves on to group 5 and so on.

Although not a different concept, it's worth mentioning how the algorithm handles overhead pressure whose preference is not to move (Min Move). This would help clarify the concept further. Recall that the priority group decides whether the preferred direction for that MV is achieved first. So in this case, even though overhead pressures preference is not to move, the optimizer will move it to achieve preferred optimization direction of higher priority MVs (Jet, Light Diesel etc.). The solver, however, will not move overhead pressure (priority 5) to optimize any MVs in priority groups 6 or 7. To put it in other words, only move overhead pressure to maximize Heavy and Light Diesel, but do not move it to help minimize overhead temperature.

Handling MVs with No Preference

Until this point we have not mentioned how the algorithm handles MVs with no preferences. Smart Tune configuration allows the user to not specify any preferences for some MVs. Users also have a choice to select any of these MVs to run against specific CV constraints.

In the CDU example, the top and bottoms circulation flows and temperatures do not have any optimization preference. In the Smart Tune workflow, these MVs can be configured to run against specific CV constraints. Assume that we want the top circulation flow MV to run against the high limit of a CV, for example the top circulation ratio. If we choose to do that, the optimizer will move the top circulation flow MV towards a direction that drives the top circulation ratio CV higher. This is done only if none of the MVs in the higher priority groups move against their preferred optimization directions. If after doing this there is still some degree of freedom left for the other three MVs with no preferred optimization direction, these MVs will not be moved in any direction (treated as minimum move).

Smart Tune Algorithm Summary and Workflow

In Aspen DMC3, a user simply specifies the preferences and priority of MVs. There is no LP tuning involved. Note that Smart Tune technology is not an engine to calculate LP costs that can be used in a traditional APC optimizer. It's an enhanced algorithm with multiple objective functions that enable the user to specify MV priorities and preferences explicitly. The controller objectives are achieved by using these MV priorities and preferences in a multiple objective sequential optimizer where, unlike a traditional APC optimizer, the magnitude of LP costs do not play a key role.

Below is a simplified version of the Smart Tunes algorithm that summarizes how the new optimizer works in achieving the controller objectives using the CDU example introduced earlier.

1. Start by optimizing the first priority group. Maximize feed pass 1 & 2, and assume that the rest of the MVs are free (no cost) to move.
2. Maximize furnace COT by keeping the same feed pass flows and assume every MV in lower priority groups (Jet flow and onwards) are free to move.
3. Use LP costs to trade-off only MVs that are in the same priority group. These LP costs do not affect MVs in other priority groups.
4. Continue until all priority groups are solved sequentially.
5. MVs with no preferred optimization direction that are assigned specific CV constraints are moved in the direction that pushes that CV towards the specified constraint.
6. If there is still some degree of freedom left, then treat the rest of the MVs (with no preference) as a minimum move.

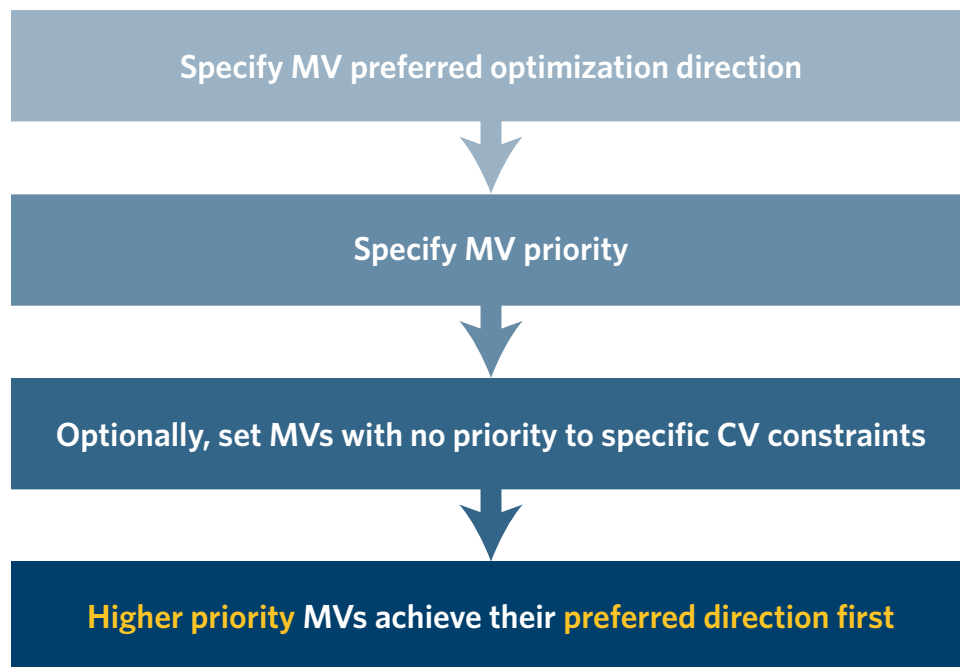


Figure 6: Smart Tune Workflow

Are LP costs still available? Yes, LP costs are still available for the user but the magnitude of these costs have little meaning. They obviously don't mean the same as they would in Aspen DMCplus. The reason for this is because the LP costs are relevant only for MVs that are in the same priority group (one of the many objective functions in Smart Tune) and do not mean much to other priority groups. Take for example the CDU priority group 3 which has only the Jet Flow MV. Jet Flow is set to maximize, so the LP cost will be negative. Now whether we have a cost of -1 or -1000 for Jet Flow, the Aspen DMC3 solution will not change. In both cases, Jet Flow will be maximized over Light and Heavy Diesel (lower priority group) but not over Feed Pass 1 & 2 (higher priority group). We know that having a Jet Flow MV cost of -1 vs -1000 will drastically change an Aspen DMCplus solution where Jet Flow will be competing with Feed pass flows since all MVs are in the same objective function. Also note that the LP costs in Smart Tune are only used to trade-off MVs that are in the same priority group.

Advantages of Smart Tune

The biggest advantage of this new technology is that it eliminates the need to perform LP cost tuning while commissioning a new controller or revamping/maintaining an existing one. This drastically reduces the efforts and complexity involved in designing and maintaining APC controllers. There are several other benefits achieved due to the way the algorithm solves.

Easily Change Controller Objectives

With previous APC technologies, changing the controller objectives can be a daunting task for many, especially if the size of the controller is medium-sized or larger. With Smart Tune, the user can simply change the priorities and achieve the intended objectives in minutes without going into any tuning, simulation or other LP tuning workflows. This is because the MV priorities and preferences are directly specified in the controller instead of MV costs. Take for example the CDU APC controller. If the economics or operational constraints now dictate that we maximize Heavy and Light Diesel over Jet Flow, we would simply change Heavy and Light Diesel priority to 3 and Jet Flow to 4 to achieve the intended objectives. With a traditional controller, doing this change requires carefully changing LP tuning so as to not affect the rest of the controller strategy.

Model Updates Do Not Affect Optimization Strategy

Recall that with traditional optimizers the steady-state solution is sensitive to model gains because all MVs are in the same objective function and the optimal solution is affected by a combination of LP costs, model gains and MV/CV constraints. When models are updated, LP tuning must be verified again to make sure the controller objectives are not affected. This issue is eliminated with the use of Smart Tune because now each priority group has its own objective function. This means that the solution is not sensitive to LP costs or model gain changes as long as the sign of either does not change.

Simplified Workflow

Smart Tune algorithm comes with an intuitive workflow and a user interface that allows the user to explicitly specify MVs preferred optimization direction and priority. An online strategy view provides an easy way to quickly view the order of MV movements for its associated CV constraints. This online strategy web view can also be used by operators and APC engineers alike to change priorities, troubleshoot, or quickly identify the controller strategy.

Conclusion

For quite some time, APC tuning and configuration has demanded a very specialized skill set. The new Sequential Multi-Objective Optimization (SMO) technology, introduced in Aspen DMC3, is a step change from traditional APC tuning methodologies. It fundamentally changes the way APC controllers are tuned and configured. The new algorithm eliminates any LP tuning efforts by enabling users to explicitly specify the controller objectives, i.e., MVs preferred optimization direction and priority or preferred CV constraints. Advantages of this new technology are outlined below:

- Eliminates steady-state LP cost tuning
- LP strategy is not sensitive to model updates
- Easily modify controller objectives by simply changing the priorities, without the need to revisit any LP cost tuning
- Greater insight into controller strategy for operators and engineers
- Enables more users to adopt APC technology with ease

To learn more about how to get started with this technology, please download the jump start guide here: www.aspentech.com/Jump-Start-Smart-Tune/.

To download or upgrade to aspenONE® software that includes this new technology, please visit the following website: www.aspentech.com/products/order-now.

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