

Platform for Advanced Control and Estimation (PACE): Shell's and Yokogawa's Next Generation Advanced Process Control Technology

R. Amrit, W. Canney, P. Carrette, R. Linn, A. Martinez, A. Singh, T. Skrovanek,
J. Valiquette, J. Williamson, J. Zhou

Shell Global Solutions (US) Inc., Shell Technology Center, Houston, TX 77082-3101.

B.J. Cott

*Shell Global Solutions International BV, 1031 HW Amsterdam, NL
(Tel: +31-20-630-3337; e-mail: b.cott@shell.com).*

Abstract: Every ten years or so, Shell has looked to refresh its Advanced Process Control (APC) technology. The last major technology upgrade occurred in 2003 when Shell along, with our APC alliance partner, Yokogawa, released SMOPro (MPC) and RQPro (quality estimation). In 2011, Shell and Yokogawa agreed to initiate the development of our next-generation APC technology. Brought to the market in 2015, the Platform for Advanced Control and Estimation (PACE) was built from the ground up, leveraging our long combined experience in APC.

1. INTRODUCTION

Shell¹ has had a long history in gaining significant benefits by the use of advanced process control (Cott 2007, Cott 2008, Cott 2012). Advanced process control allows us to manage production, energy efficiency, product quality and waste minimization in a coordinated and consistent manner.

SMOPro, our current technology released to the market in 2003, was a typical design of its time: a two-level linear controller with a static optimization layer to compute the targets for the lower level dynamic control problem.

Our next generation technology, the Platform for Advanced Control and Estimation (PACE) represents a quantum leap that will simultaneously drive down the costs of developing and maintaining APC applications while providing new tools and capabilities to address even more challenging control problems and deliver more value to our customers. While PACE is still based on a two-level controller design, many of the assumptions used in SMOPro have been further relaxed.

In this paper, we will discuss the major advances made in various aspects of the technology, including how we formulate and solve the estimation and control problems and how we minimize the impact of maintenance activities on the running controller. Finally, we will discuss the challenges and opportunities for further work in this space.

¹ The companies in which Royal Dutch Shell plc directly and indirectly owns investments are separate entities. In this publication the expressions "Shell", "Group" and "Shell Group" are sometimes used for convenience where references are made to Group companies in general. Likewise, the words "we", "us" and "our" are also used to refer to Group companies in general or those who work for them. These expressions are also used where there is no purpose in identifying specific companies.

2. ESTIMATION

Traditionally, the estimation part of many control algorithms has taken a secondary role to the control law in the development of the algorithm. PACE continues our tradition of focusing as much attention on the estimation problem as the control law. Our experience has shown that the better we can predict where the process is going, the better the overall control performance will be. Fundamentally, the estimation problem is one of using process feedback to manage the presence of various deterministic and stochastic disturbances that always presented during real operation (Muskie and Badgwell (2002), Pannocchia and Rawlings (2003), Rajameni et al (2009)). These include:

- Plant model mismatch & model uncertainty;
- Unmodeled / unanticipated perturbations;
- Sensor / measurement noise.

PACE and its predecessor SMOPro (Cott 2007) achieve better performance in the estimation space by:

1. Allowing the use of additional measurements (called process output variables (POV's)) in the estimation problem that do not necessarily participate in the control law but improve the performance of the estimation problem;
2. Relaxing the typical assumption that the structure between manipulated variables (MV's) and process out variables (POV's) must be a dynamic matrix;
3. Allowing the designer to match the prediction error update function for each process output variable to the behaviour of that process output variable.

A control designer sitting down with the PACE software begins by building the estimator, which relates MV's and DV's (disturbance variables) to POV's. POV's can also feed

other POV's which provides more much flexibility in defining the relationship amongst variables, in particular, around how prediction errors are forecasted to other POV's Only later in the design process does the control designer select POV's to become controlled variables (CV's) in the control law. Model identification using POV's is often simpler and more robust as the control designer is fitting a series of models between POV's rather than a MV-to-CV relationship which may have many process operations between them. PACE convolutes the POV-to-POV models internally to arrive the equivalent control law.

PACE further relaxes model structure assumptions by allowing the control designer to use a different model structure at the control layer (Figure 1). While the default control layer shares the same model relationships, there are times when changing the model structure for control proposes can be very beneficial. For example, there are often good reasons to segregate manipulated variables so that they only participate on a subset of controlled variables, even though they have strong impacts throughout the controller. Keeping the full relationship in the estimation layer, while removing models at the control layer, can deliver the desired controller behaviour.

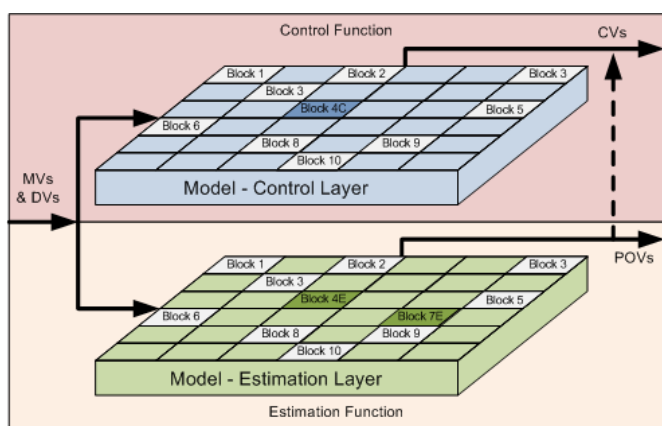


Figure 1 - Control and Estimation Layers

One key set of POV's that SMOCPro and PACE leverage heavily in the estimation problem is the base layer controllers that they write to as manipulated variables. While it is possible to assume that the base layer controller will always drive the process output variable to the setpoint requested by PACE, it has been our experience that this assumption does not hold true very often and modelling the performance of the base layer controller is very important for the estimation problem.

Properly modelling the behaviour of these base layer control loops to include the effect of the base layer control loops significantly quiets the control actions computed by a controller. In one example, incorporating the full behaviour of a furnace outlet temperature in the estimation problem reduced the amount of control action on a crude column by about 50%. This indicated that the original control design was over-controlling the column, taking action for errors seen in the controlled variables that were in fact already corrected for in the base layer control.

While SMOCPro provided a straightforward way to model the relationship between the base layer controller setpoint and the actual process output variable under control, PACE extends this to also model the behaviour of the control valve which allows the control designer to much more precisely manage saturation conditions. There are many times when the designer does not want the base layer control to saturate, but in a few cases saturation assures that the controller is making maximum use of the variables it has available. This also allows the control designer to better handle changes in controller modes. For example, if PACE was manipulating a flow to control a level and later the flow was taken out of remote and returned to local, PACE will properly model this change in behaviour of the flow controller. This was not possible with SMOCPro.

Another innovation relates to soft sensor technology. In our previous technology, the implementation of soft sensors (also often referred to as quality estimators) was done in a separate software package (RQEPro) and we found that this often caused us to repeat a great deal of information in both RQEPro and SMOCPro. Fundamentally the major difference between the two packages was only in the algorithm to do the prediction error updates. In PACE, the control designer has the ability to select which error prediction update method to use: either the PACE standard method (used for measurements scanned at a frequency equal to or fast than the controller or an error projection update routine based on our soft sensor technology for infrequently sampled measurements. This second update method allows the control designer to bring lab data and slow analyser feedback directly into the estimation layer, eliminating the need to bring a separate application

Given the wider number of options available to the control designer during construction of the estimation layer, it is vitally important to provide powerful visualization tools to display to the designer the current model construction. PACE permits a variety of visualizations of the estimation layer (and control layer as well) including an extended dynamic matrix structure (Figure 2) that shows POV's in both the rows and columns of the dynamic matrix. Convolution of the individual models into the final dynamic models is performed automatically by PACE.

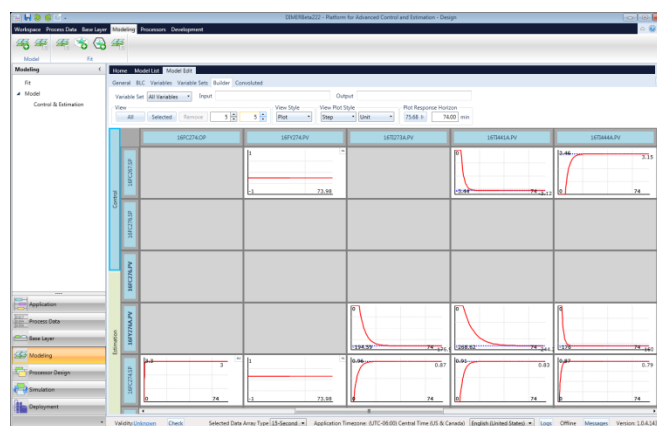


Figure 2 - Extended Dynamic Matrix Representation

PACE also provides a graphic model viewer (GMV) that shows the cause-effect structures amongst the MV's, DV's, POV's and CV's (Figure 3):

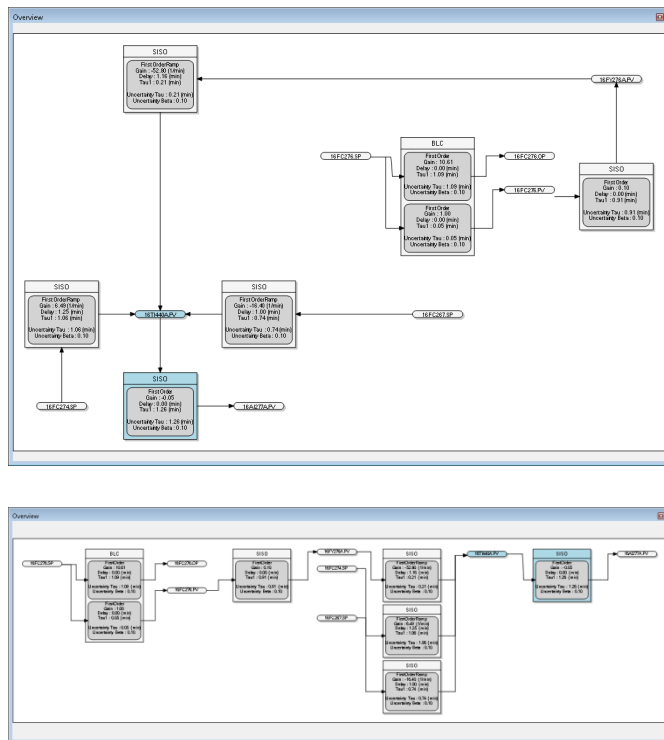


Figure 3 - Graphic Model Views

The advantage of the cause-effect graphic model view is that it shows the control designer which POV's impact which CV's and if particular CV's are highly correlated with each other.

With all of this focus on POV's, a good rule of thumb we use is that a well-designed estimation layer has at least two POV's for every CV. Since at least one POV is directly tied to each CV and each MV has its current value included in the estimation model, this rule of thumb is relatively easy to achieve for refining and chemical processes by adding only a few more key POV's: as an example, for distillation columns, these additional POV's tend to be secondary tray temperatures and/or column loading indicators like delta pressure measurements.

Finally, as the span of control strategies continues to increase, it is more and more likely that the estimation layer will have to deal with parts of the plant being taken out of service for short-term maintenance. A good example is the shutdown of an ethylene cracker furnace for decoking. It is very important that the estimation layer properly handling this shutdown state. To help with this, we have included an equipment out-of-service concept that allows the control designer to take groups of CV's, POV's, MV's and/or DV's out of the estimation and control layers in a structured manner.

3. CONTROL AND OPTIMIZATION

While we have focused a great deal of innovation into the estimation layer of PACE, the control layer was also fully revised, starting with economic functions. We took the opportunity to move from a traditional Quadratic Program solver technology and embraced a state-of-the-art solver technology that permits the use of general nonlinear functions for economic functions, not just linear or quadratic functions. This creates the opportunity to design and implement economic functions that align with the business drivers rather than relying on a local and simplified approximation.

Furthermore, PACE permits multiple economic functions that can be optimized simultaneously. The control designer can prioritize economic functions against each other and versus CV constraints (Figure 4). These economic functions are optimized in the order of decreasing priorities until all degrees of freedom are exhausted. If specified as equal priority, the control engineer can weigh the economic functions against each other.

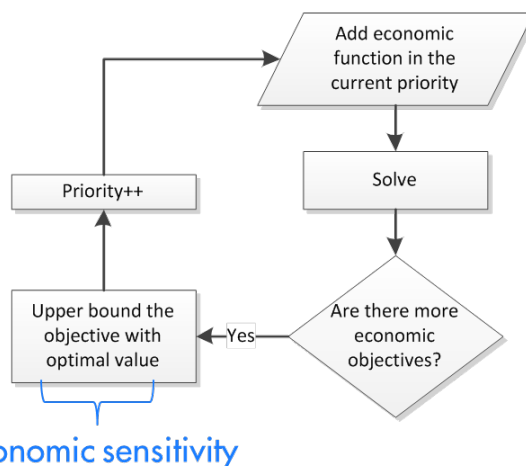


Figure 4 - Economic Optimization in PACE

Moving from the steady-state optimization space to the dynamic control problem, we have moved away from SMOCP's limited horizon formulation to an infinite horizon formulation in PACE. To do this, we switched from having the system attracted to its steady state target using MV clamping in SMOCP to having the system attracted to its steady state using penalties on CV deviations. This allows a longer horizon for the MV's to plan moves.

We discovered with SMOCP that relying on the same set of tuning parameters for both the static optimization and the dynamic control problem caused challenges for the control designer to consistently get the desired behaviour, PACE splits the tuning for the static and dynamic control into different parameter sets, resulting in a much more straightforward way to tune the dynamic response.

Another aspect of the dynamic control problem in SMOCP that we wished to improve was the handling of unstable or ramp processes. Unstable or ramp behaviour in advanced process control design often come from the incorporation of level control of tanks and/or accumulators, but there are also several refining and chemical processes that exhibit unstable

behaviour (for example, temperature control in partial combustion catalytic crackers). SMOCPRO exerted very tight control on ramp imbalances which often unnecessarily constrained the control problem. PACE allows a temporary relaxation of the ramp imbalance constraints when within limits. This is very beneficial when controlling large tanks with a slow dynamic cycle. (e.g. withdrawing inventory during the day and replenishing it at night). This also helps to make tuning of levels and ramps more straightforward.

Finally, we have investigated how to best deal with the handshake between the steady state optimization problem and the dynamic control problem. This has remained a major challenge when integrating APC with nonlinear real-time optimization (RTO) systems, where the optimal operating point is computed as a set of coordinated RTO targets.

Traditionally, the handshake between APC and RTO assumes that each RTO target is independent of each other and a simple weighting of errors between the RTO targets and the equivalent CV's in the dynamic control problem is sufficient. In many cases, the formulation of the RTO system results in the economic function contours to be functions of several RTO targets. Failing to incorporate these higher-order relationships into the computation of the targets for the dynamic control problem will result in the controller not fully extracting the full cumulative benefits available to it as discussed by Rawlings et al (2008) and Rawlings and Amrit (2009).

For PACE, the Best Performance Value (BPV) concept bridges this gap. The BPV is defined as an optimized set of CV values that correspond to the highest economic function value but lie within the CV space limits (). The static target that the dynamic control problem uses is easily computed alongside the BPV.




Figure 5 - Best Performance Value Concept

The BPV concept therefore ensures that the relationship amongst the CV's that generates the highest economic values are preserved upon the calculation of the static targets and the controller will be able to achieve a good portion of the cumulative benefits available.

4. CONTROLLER MAINTENANCE

While the previous two sections have focused on the algorithmic aspects of PACE, there are many equally importance software design aspects that significantly increase the productivity of the control designer in delivering and maintaining APC strategies.

Changing a controller design while it is already active has traditionally been a challenging problem. In the past, with restricted memory and CPU space, there was little more that could be done than to stop the old controller running, upload the new control design and then restart the new controller. This created a downtime in the controller operation and perhaps more subtly it also took a long period of time for the controller to rebuild its predictions from this restart point. Fundamentally, this created the opportunity for a significant period of poor performance whenever a controller was updated.

In PACE, we allow both staged and live controllers (Figure 6). Configurations can be archived, both within the design-time and run-time environments. In the run-time environment, controllers can also be staged which means they are running in estimation mode alongside the active controller or even pull historical data from the running controller to build predictions. The predictions on the staged controller are therefore much closer to those of the active controller, so when the staged controller is promoted to being the active controller, the controller effectively does not go offline nor does the controller performance degrade as the predictions are in good shape.

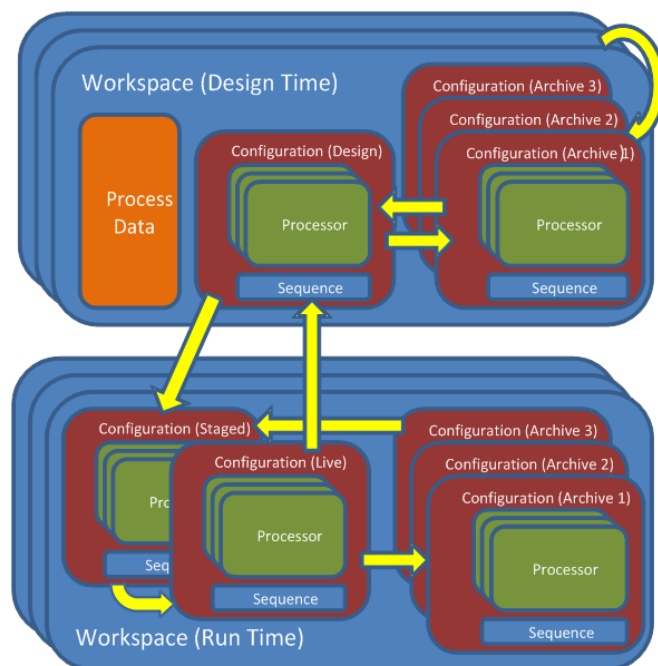


Figure 6 - Design and Run Time Staging

Another reason why APC strategies have traditionally gone offline is for the control designer to change or update a model between two variables in the controller. PACE permits a great deal of flexibility in changing the individual models in the control or estimation layer. Both the transfer function type (first order, second order, ramp) and the transfer function parameter list (gain, delay, time constants) can be changed online without putting the controller in standby. While the control designer can define multiple elements for a single input/output pair, only one element can be active per SISO relationship. PACE allows the user to programmatically switch amongst these multiple elements to account for operation mode changes.

5. CHALLENGES AND OPPORTUNITIES

With the size of controllers ever increasing, it is getting more and more difficult for the control designer to confirm whether the controller is, in fact, performing as expected.

At the simplest level, this type of question is often posed as: "Looking back to a moment in time, why did the controller do what it did?" This is a challenging question as it requires not only deep knowledge of the controller design and implementation, but also the history of the data presented to and processed through the controller. This moves the diagnosis of controller behaviour into the "big data" realm. Of course, this approach only works if the future behaviour is similar to that of the past. While this assumption often holds, there are many situations where a change of process behaviour due to a physical change in the process (damaged processing equipment) or operating goal (new product specification) makes fully relying on past data impossible.

More generally, even if there may not be an event that initiates investigation, the control designer should be spending some time asking and answering the following questions:

- Could the controller be generating more benefits than it currently is?
- Are we making the best use of the full operating window?
- What are the trade-offs to be made in order to grasp those opportunities?

To answer these questions, it is clear that the economic landscape of the process operation must be known to a high degree of precision and that it must be overlaid onto the controller's performance. It also indicates that some degree of experimentation needs to be included in order to search for new solutions that may not be represented in the current controller design.

SUMMARY

Getting the opportunity to write from the ground up a new advanced process control technology allowed Shell and Yokogawa to revisit many of the design decisions we made in our previous technology. What we found was that many of

the design decisions were related to the fundamental constraints of software and computing technology of ten or more years ago.

With access to today's modern computers and software languages, we have been able to advance many of the areas within advanced process control technology to allow the delivery of lower-maintenance, higher-performing controllers within a shorter project timeline.

REFERENCES

- Cott, B.J. (2007). "Unit-Wide Model Predictive Control with SMOCP", D.G. Fisher Award Keynote Lecture, CSChE Annual Conference, Edmonton, Canada.
- Cott, B.J. (2008) "Stretching the boundaries: Expanding the use of APC into non-traditional areas", Plenary Lecture, ADCONIP 2008, Jasper, Canada.
- Cott, B.J. (2012). "Extracting Maximum Economic Benefits from Nonlinearities in the Refining and Chemical Industries", Plenary Lecture, IFAC Conference on Nonlinear Model Predictive Control 2012 (NMPC'12), Noordwijkhout, Netherlands.
- Muske, K. and T. Badgwell (2002), "Disturbance modeling for offset-free linear model predictive control," *Journal of Process Control*, vol. 12, no. 5, pp. 617-632, 2002.
- Pannocchia, G. and J. B. Rawlings, (2003) "Disturbance models for offset-free model-predictive control," *AIChE Journal*, vol. 49, no. 2, pp. 426-437, 2003.
- Rajamani, M., J. Rawlings and S. Qin, (2009) "Achieving state estimation equivalence for misassigned disturbances in offset-free model predictive control," *AIChE Journal*, vol. 55, no. 2, pp. 396--407, 2009.
- Rawlings, J.B., D. Bonn , J. B. Jorgensen, A. N. Venkat and S. B. Jorgensen (2008) , "Unreachable setpoints in model predictive control," *IEEE Transactions on Automatic Control*, vol. 53, no. 9, pp. 2209-2215, 2008.
- Rawlings, J.B., and R. Amrit (2009), "Optimizing process economic performance using model predictive control," in *Nonlinear Model Predictive Control*, pp. 119-138, 2009.