

Implementation and Rationale for PJM's Conditional Neutrality Regulation Signals

PJM Staff
January 2017



Contents

Acronyms	3
Introduction	4
Balancing the Demand for Energy	4
Frequency Regulation Ancillary Service	5
Previous Controller Design	6
Conditional Neutrality Controller Design	7
Design Philosophy	7
Implementation	8
Implications	10
Conditional Neutrality Example	11
Simulation and Response Models	15
Seasonal Analysis for Regulation Requirement	18
Regulation Market Requirement	20
Scenario Generation	20
Rate of Technical Substitution Curve Formation	22
Measurement and Verification of Substitution Curves	25

Acronyms

ACE	Area Control Error
ACS	ACE Correction Signal
AGC	Automatic Generation Control
AREG	Assigned Regulation
BAAL	Balancing Authority Area Control Error Limit
CPS	Control Performance Standard
EMS	Energy Management System
FERC	Federal Energy Regulatory Commission
LOC	Lost Opportunity Cost
NERC	North-America Electric Reliability Corporation
PBR	Performance Based Regulation
PI	Proportion-Integral (controller)
PID	Proportion-Integral-Derivative (controller)
Regulation A	Traditional Regulation (signal type)
Regulation D	Dynamic Regulation (signal type)
RTO	Regional Transmission Organization
RTSCED	Real-Time Security Constrained Economic Dispatch
TREG	Total Regulation Capability

Introduction

On Jan. 9, 2017, PJM implemented new regulation signals and requirements for regulation service. These changes were driven through the PJM Stakeholder Process, being initiated by a problem statement from PJM and the Market Monitor in October 2015 describing the observed Regulation D signal not always providing ACE control and the potential reliability impact.

The PJM Regulation conditional neutrality controller, put into production on Jan. 9, 2017, was engineered to provide a balance between RTO ACE control and energy storage neutrality needs. Both Regulation A and Regulation D signals now move together in the direction that best minimizes RTO ACE, correcting the problem identified with the original Regulation D signal formation. The controller introduced a new neutrality bias to the Regulation A signal, conditional on having additional capability between optimal correction and total regulation capability limits. The neutrality bias intentionally over-corrects with the Regulation A signal, such that Regulation D can move in the direction that reduces accumulated state of charge. The result of the two controlling actions, when taken together in concert, is regulation which meets the short time horizon needs of RTO ACE control; while, fulfilling the longer-time-horizon needs of energy storage generation and ramp limitation of traditional resources.

Balancing the Demand for Energy

The primary role of a Balancing Authority is to manage the supply and demand of electricity. Companies such as PJM serve this role by economically dispatching generation to meet real-time load and interchange on the bulk electric system. However, changes in supply and demand are not precisely predictable and there will be mismatches resulting in non-zero ACE. Regulation is an essential reliability product that PJM relies on to help manage ACE.

Regulation Services aid in the continuous balancing of generation and load, assisting in maintaining interconnection frequency and Balancing Authority ACE. Regulation is a variable amount of generation under automatic generation control (AGC) that PJM sends a regulation signal to raise or lower the resource's output to correct for instantaneous changes in load and generation.

As a measure of control, PJM utilized the RTO ACE control required by the NERC BAL-001 "Real Power Balancing Control Performance" standard. There are two required measurements of control performance, the Control Performance Standard 1 (CPS-1) and the Balancing Authority ACE Limit (BAAL). Both of these are measured by comparing the RTO ACE against measured frequency deviation over periods of time. Figure 1, PJM Control Performance, displays the monthly average performance of CPS-1 and BAAL over 13 months. In April 2016, PJM, in accordance with NERC BAL-003-1 requirements, implemented a 15 percent reduction in frequency bias which measurably affected the CPS-1 calculations, reducing average CPS-1 score by 13 points month to month.

Additionally, on July 1, 2016, PJM updated the CPS-1 calculation to utilize reporting ACE in replace of Control ACE, in accordance with the NERC BAL-001-1 effective date. This resulted in another measurable impact to the CPS-1 scores. Despite these changes, the average CPS-1 score remained above 100 (as designated by the dashed line) as required by the CPS-1 standard.

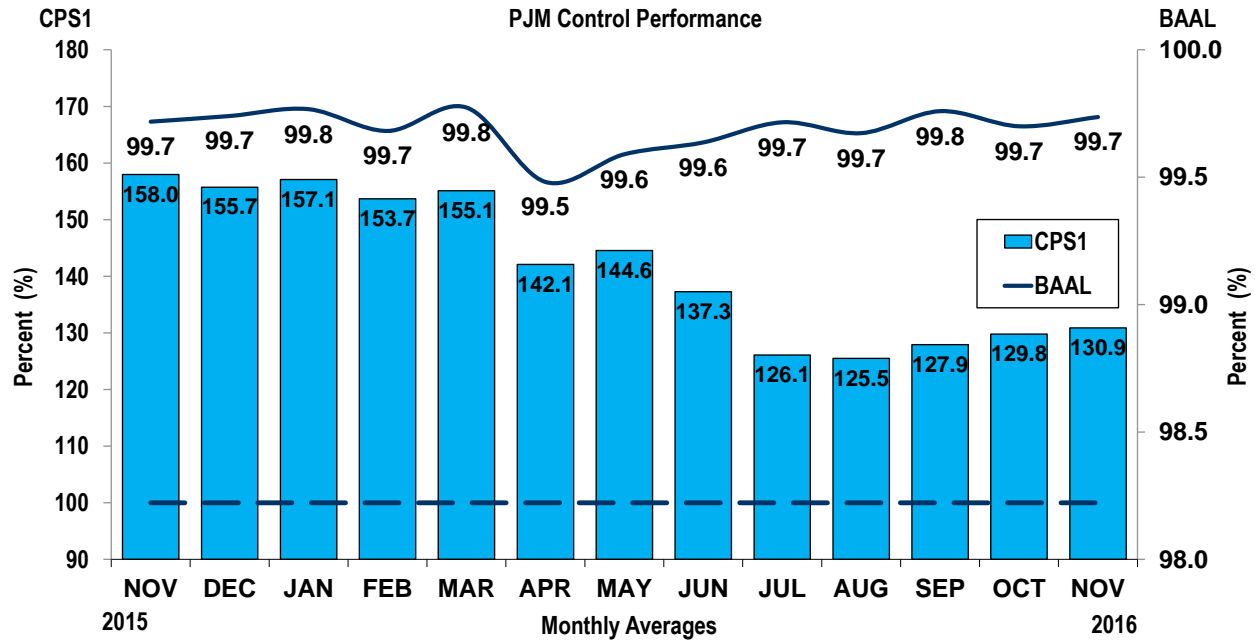


Figure 1: PJM Control Performance

Frequency Regulation Ancillary Service

Since 2002, PJM operated a Regulation Market to competitively assign regulation ancillary service obligations, in least cost merit order, to meet a total regulation requirement. In 2012, PJM redesigned the Regulation Market to meet FERC Order 755, and the design is now referred to as Performance-based Regulation (PBR). In the Performance-based Regulation design, regulating resources offer, clear, and settle based on a measured performance of response to assigned regulation signal.

As of 2012, PJM generates two types of regulation signals: the traditional Regulation A signal that is designed to accommodate traditional ramp-limited resources, and the dynamic Regulation D signal that is designed to accommodate faster-moving energy-limited resources. PJM manages the total amount of regulation on the system at any given time by a reported Total Regulation Capability (TREG) value. TREG is reported as a fleet aggregation of resource capability, which represents the total amount of movement a resource, or group of resources, can move up and down relative to a basepoint. PJM operates a bi-directional, symmetric regulation signal for both regulation signals, such that the signal sent to the resource can vary from positive TREG (a full raise signal) through the entire range down to negative TREG (a full lower signal), centered on the set basepoint for the resources. The EMS's AGC generates new regulation signals with a calculation frequency of two seconds and sends them to the fleets via communication protocols.

Today, PJM has three major groups of resources participating in regulation service: traditional generation (e.g., steam, hydro, CT, etc.), energy storage (e.g. batteries, flywheels, etc.), and demand response. For a traditional resource, a regulation signal directs the unit to increase or decrease its energy output relative to a regulation basepoint. Because these regulating resources also participate in the energy markets, the regulation basepoint is

often the RTSCED-generated economic basepoint. Additionally, units that are providing Regulation A service are typically ramp-limited; the RTSCED operation reduces the ramp-rate used in economic basepoint changes by a calculated amount proportional to the regulation assignment. This calculation takes place to honor the energy bid-in ramp-rate of the regulation resources.

For an energy storage resource, a positive regulation signal represents a request to inject energy into the grid, and a negative signal represents a request to withdraw energy from the grid. Unlike a traditional regulating resource which burns fuel to generate, an energy storage resource will eventually deplete its stored charge if it were to continuously inject energy into the grid. Energy storage resources have the ability to supply PJM with a load basepoint in order to manage their state of charge. This load basepoint reflects the energy withdraw or injection of the resource outside of following the regulation signal. Due to the energy-limited nature of storage resources, the Regulation D signal attempts to keep the signal neutral over time when the total regulation product has the capability to do so. This neutrality objective is intended to assist with any energy bias seen through these signals.

Demand response is the final resource group which typically participates in the Regulation Market. These resources follow the regulation signal by adjusting their withdraw of energy from the grid. As demand response resources cannot inject energy to the grid, these resources utilize a load basepoint to center the resources at a point they can regulate around by increasing or decreasing their energy withdraw from the grid.

Previous Controller Design

The previous PJM Regulation Controller was less effective in optimizing ACE control than the Conditional Neutrality Controller, as identified in a few areas. First, the Regulation A and Regulation D signals were created independent of each other, not allowing for coordinated ACE control between all regulation resources. Second, the Regulation A acceleration function present in the controller provided signals that could exceed resource expectations. Third, and most measurable, the “hard-neutrality” function of the Regulation D signal sent a measurable amount of resources, at times, in the opposite direction of ACE control.

From the inception of PBR in October 2012 through Jan. 9, 2017, PJM's AGC regulation signals were generated from a modified proportional-integral controller (PI controller) design. The controller begins with an ACE Correcting Signal (ACS, equaling negative ACE), then applies a small dead-band reduction before sending it into an integrator, which is a modification of a classic PI controller. The PI controller design is useful for managing both rapid fluctuations and steady state responses to an input signal. In this controller model, only the traditional Regulation A signal is dependent on the PI controller output, while the Regulation D signal is independently generated from a proportional scalar of the ACS.

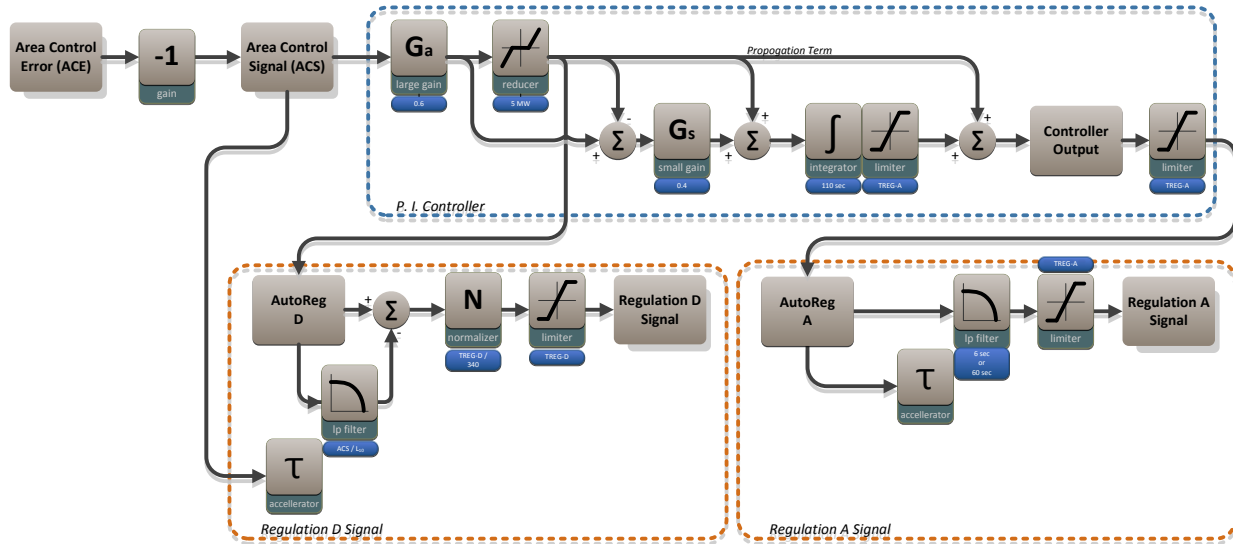


Figure 2: AGC Controller Prior to January 9th 2016

The PI controller output is directed to a dynamic low-pass filter where it becomes the Regulation A signal. The time-constant of this low-pass filter is a function of the sign of the filtered signal relative to the sign of the PI controller; if the two are opposite signed, the time constant is shortened, resulting in an acceleration effect that improves the convergence of the filtered signal back to the PI controller signal. The acceleration of the Regulation A signal could result in a signal that changes faster than the expected capability of the participating resources.

Next, the proportional term of the PI controller is directed through a separate dynamic low-pass filter into the Regulation D formation path. The time constant of the second filter is a function of the ACE deviation, where a larger deviation becomes a larger time constant. The Regulation D signal is determined as the residual of the proportional term and the second filter, which results in a signal that exhibits zero-centered “neutrality.” Large ACE deviations pass through the proportional term of the PI controller, where they result in larger short-term swings in the Regulation D signal. Over time, the filter catches up to the deviation, and the residual converges the Regulation D signal back to zero (the appearance of neutrality). The Regulation D signal is not explicitly neutral, in that it does not balance the accumulation of signal deviation, and exhibits a slight bias. Under this controller design, Regulation D resources may adjust their output to maintain state of charge, by altering their regulating set points to offset the signal bias or dropping out of the regulation service to rebalance their state of charge.

Conditional Neutrality Controller Design

Beginning on Jan. 9, 2017, PJM began controlling regulation resources using the new conditional neutrality controller as described in the following sections.

Design Philosophy

The conditional neutrality controller is a hybrid proportional-integral-derivative controller (PID controller) model with an internal feedback loop which controls for the state of charge of Regulation D resources. The controller retains the

two signal approach to ACE control, with the residual between the ideal control (that is, ACS) and the traditional signal being sent to the faster resources.

Unlike the previous controller, the conditional neutrality controller does not forcibly converge the Regulation D signal back to zero within a fixed period of time. Instead, Regulation D continues to provide the residual signal that would best control RTO ACE to zero. In this way, Regulation D is providing short-term ACE control faster than the traditional Regulation A, with energy to be later balanced by the Regulation A resources when available to do so.

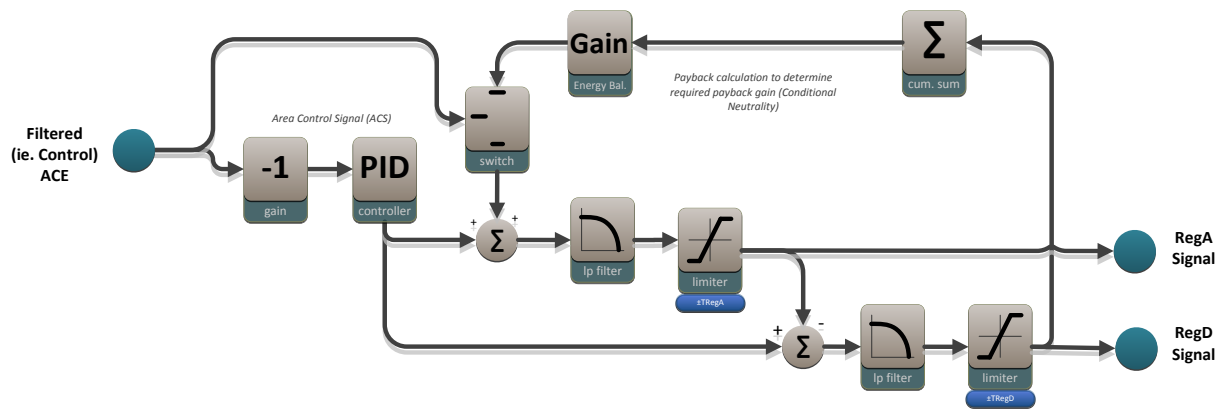


Figure 3: Conditional Neutrality Controller, High-level

As the Regulation D signal is followed by the energy storage resources, the controller measures the accumulation of the instantaneous injection and withdraw (measured in megawatts (MW) in discrete two-second intervals), which equates to total megawatt-hour (MWh) depleted and charged from the resource. After the ACE deviation is stabilized through normal optimal control, if there is additional available regulation capability to move the Regulation A resources, then the Regulation A resources are additionally moved by a function of the Regulation D energy accumulation. This has the net effect of “over-steering” with the traditional Regulation A resources, which causes the Regulation D resources to move in a direction that reverses the state of charge accumulation. The conditional neutrality also continually checks the current position of ACE when determining the payback. If the ACE deviation is very large, the controller places extra priority on ACE control by turning payback off until the regulation resources again have ability to assist with energy management.

Implementation

This controller design also begins with an ACE Correcting Signal (ACS, equaling negative ACE), then sends this into a PID controller. The PID controller output plus the neutrality bias is sent to a low-pass filter, where it becomes the Regulation A signal. The residual between the Regulation A signal and the controller output is sent to a separate low-pass filter, which becomes the Regulation D signal. In this design, the summation of the regulation signals always provide the optimum ACE control determined by the PI controller output.

Taking the implementation described thus far, the controller would coordinate the traditional Regulation A resources and dynamic Regulation D into a signal that is effective at correcting ACE deviations. However, because Regulation A resources cannot always support the energy needs of Regulation D resources, the ability of these Regulation D

resources to follow the signal could diminish over extended periods of time, causing the response to be less effective at ACE control. Accommodation of the energy capability limitations of Regulation D resources is analogous to the accommodation of ramp capability in a traditional AGC formulation; this is necessary for effective control to ensure response meets control objectives. To implement the Regulation A signal, the controller includes a low-pass filter with a large time constant to address this ramp capability restriction. An analogous dynamic functionality is also included to meet energy capability needs of Regulation D resources: this is the conditional neutrality function of the controller.

The conditional neutrality portion of the controller estimates a total charge capacity, as a function of the modeled energy storage size and the reported total regulation capability. Resources are assumed to operate around a 50 percent state of charge starting point, so that the devices avoid extreme charge or discharge states. This creates an effective limit to the stored energy as half of the reported TREG capability, multiplied by the reported energy storage duration.

The Regulation D signal sent to the resource is integrated to capture the net injection or withdraw of energy over time. Based on this accumulation, the controller estimates where the state of charge should be as a percent of the storage limit, which is represented as 100 percent fully charged to 0 percent at fully depleted.

An example of the state of charge calculation is displayed in Figure 4 below. This example shows three different battery resources starting at a half charged state at time zero. Then the modeled batteries are asked to provide a full lower in the Regulation Market which results in an increase to state of charge, until the battery resources are at a fully charged state. The Figure 4 graphic demonstrates the energy capability of different batteries based on their storage duration. To continue the example, a 15-minute battery starting at 50 percent state of charge and asked to charging at its maximum energy withdraw will reach a 100 percent state of charge in 7.5 minutes. Furthermore, the same resource could reduce its output to 50 percent of name-plate capacity and be charged to 100 percent capacity in 15 minutes.

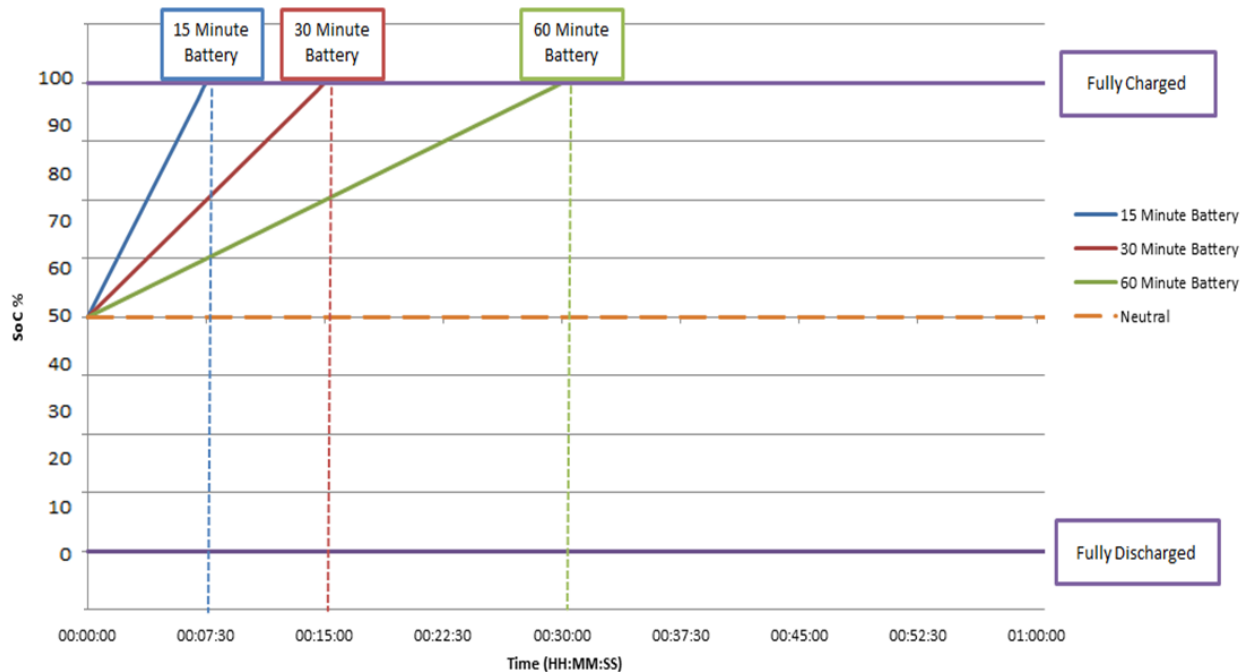


Figure 4: Energy Storage State of Charge When Sent a “Full Lower” Regulation Signal

As these resources deplete their state of charge, the controller’s variable payback gain increases in size, changing with respect to absolute percent of the storage limit. The controller calculates the Regulation A neutrality bias as the Regulation D accumulation multiplied by the variable gain. Thus, as the battery approaches a fully charged or fully discharged state, the neutrality bias increases, which further drives the Regulation A resources beyond the normal ideal control point. This action directs the Regulation A in an over-controlled direction, which means that as the resource begins to respond, the Regulation D signal will move in the direction that reduces the charge accumulation allowing them to balance out their state of charge.

Implications

It is important to note that all regulation signals remain bounded by the reported total regulation capability (TREG) values, even with the neutrality bias component. In a situation with an extreme ACE deviation, for example a high and sustained step value, the Regulation D resources will be the first to respond down to absorb energy from the grid in a sharp, decisive motion. The Regulation A signal is slowed by the filter, but also moves in the down direction with some delay in a smoothed motion. Over time, the Regulation D signal accumulator sends an additional bias to lower Regulation A, but the combination remains limited by the low-pass filter. Eventually, both Regulation A and Regulation D could become saturated at a full lower position. Regulation D will continue to receive a lower signal until the ACE deviation is arrested. In this way, the neutrality for energy storage resources is conditional upon the ability for the regulation service to control for total ACE, and neutrality is sacrificed while ACE correction is necessary. Note that if the ACE deviation is large, the controller will suspend the payback logic until such time as the ACE deviation is back within acceptable bounds; this allows the Regulation A signal to turn around even faster when ACE correction is necessary.

After some time, some of the Regulation D resources could reach their energy capability and decrease output to 0 MW because they have fully charged. Regardless, the controller will continue to send a full lower signal. Importantly; however, the integrator within the conditional neutrality calculation will no longer continue to accumulate because the resources are no longer expending energy. In the sense that the resources can no longer charge, there is also no need to even further over-steer Regulation A to discharge this energy later, once ACE has recovered¹.

After ACE is balanced, the controller will reset, and Regulation D resources will be the first to respond in the upward direction, relieving some of the accumulation. Because of the large historical accumulation, the Regulation A resources will continue to be biased down such that Regulation D resources will receive a larger residual, which will relieve the accumulation faster. After the accumulation is reduced, the neutrality bias will lower, and Regulation A will be directed back toward only the ACE correction signal.

Conditional Neutrality Example

The following images depict the process of the conditional neutrality controller as it calculates payback in order to rebalance the accumulated energy of the Regulation D signal. This example is for illustrational purposes only. The values internal to the controller are continually optimized to ensure that PJM is receiving reliable control from the regulation resources.

In Figure 5 below, the energy accumulation starts out at approximately 20 percent in the charged direction. Because there is some imbalance of the energy accumulate, a payback term is calculated which biases Regulation A in the negative direction. This bias allows Regulation D to control for the residual which causes the Regulation D signal to be more closely centered around the 0 MW point.

¹ Note that the discussion is centered on the modeled state of charge of the Regulation D resources within the conditional neutrality component of the controller, not the actual state of charge of the Regulation D resources. In principle, the model and actual state of charge are aligned; in practice, differences occur due to different capacities and operational output. This is realized by resources dropping out before or after the controller stop accumulating energy via its integrator.




Figure 5: Conditional Neutrality Example - Payback Calculated

Due to the conditions of ACE during this period, the Regulation A signal needs to continue to hold at full lower in order to control ACE back to 0. Regulation D during this time period continues to stay on the low side in order to help with ACE control which adds more accumulated energy to the signal. As the accumulated energy passes 25 percent in the charged direction, the controller adjusts the payback amount from Regulation A in order to try and rebalance Regulation D more quickly. However, Regulation A is fully utilized controlling ACE and cannot provide any additional payback during this period, as seen in Figure 6.




Figure 6: Conditional Neutrality Example - Regulation A Fully Utilized

For the next segment of time (Figure 7), ACS moves in the positive direction shifting the need for ACE control. At this point, Regulation A remains in the full lower direction (over-steer) to allow Regulation D to go in the positive (discharge) direction. As the Regulation D signal is now in the positive direction, the accumulated energy starts to decrease and passes a set threshold for the variable payback gain which lowers the payback requested from Regulation A resources.




Figure 7: Conditional Neutrality Example - Regulation D Rebalances Through Payback

As ACE is rebalanced, the Regulation A and Regulation D signals continue to provide effective ACE control. Because the energy of the Regulation D signal is still not balanced, a payback value is continually calculated to over-steer Regulation A in a way that will help center the Regulation D signal around 0 MW. The Regulation D signal continues to center around 0 MW until the accumulated energy is balanced, at which point the payback term goes to 0 MW, as seen in the highlighted section of Figure 8 below.




Figure 8: Conditional Neutrality Example - Regulation D Continues to Rebalance Through Payback Until Corrected

The controller continues to monitor the accumulated energy of the signal on a rolling basis. Based on the degree of imbalance, the controller determines how much energy is sent to Regulation A resources to bring the Regulation D signal back to a balanced state. It is important to reiterate that in instances where Regulation A resources are fully utilized providing ACE control (Figure 6 above), Regulation A will be unable to provide energy neutrality management. Once ACE is corrected to a point where Regulation A is not fully utilized, the controller will again receive payback from Regulation A resources and rebalance the Regulation D resources. This is the key concept behind the conditional aspect of the conditional neutrality controller design. The ultimate priority of the controller is to provide ACE control, and will only provide energy neutrality management when additional capability of Regulation A resources is available to do so.

Simulation and Response Models

Frequency regulation is an important element in providing RTO ACE control, therefore any modification to the AGC controller design needs to be thoroughly tested. In order to do this, the conditional neutrality controller needed to be tested in a simulated environment. The analysis methodology was to retrieve historical data for periods of time, create an “uncorrected ACE” data set by carving out response to the historical regulation signals, then execute the

new controller model with various scenarios to find an acceptable set of tuning parameters as seen below in Figure 9.




Figure 9: High Level Simulation Loop

To simulate a closed-loop response to the new regulation signals, several response characteristics were identified that would be important to model within the simulation. After signal generation, there is a nominal network transport delay as data is sent to the plant for consumption, which was modeled as eight seconds time shift delay. Upon receiving the signal, the plant model (detailed below) processes the signal, resulting in a plant output MW injecting into the grid. The current regulating fleet is composed of four resource types: thermal, hydroelectric, energy storage and demand response.

Thermal regulating resources are primarily fossil-fueled steam and combustion turbines, which burn fuels to inject energy into the bulk electric system. A thermal resource provides regulation by injecting energy above or below the economic basepoint. It is the change in output relative to the basepoint that is considered regulation response. For the simulation's purposes, only changes in generation needed to be modeled, not the entire output stack. Thermal resources were modeled as being constrained between high and low operating limits, which in the response model was the total regulation capability. Additionally, a two MW dead-band was modeled around the desired output, as determined from historical observation. Response to the desired output was modeled as a low-pass filter, which was then constrained by a ramp rate limitation that limited the change in output per step. The thermal response models were compared against historical data from existing steam power plant output for validation, which resulted in a low-pass filter time constant of 35 seconds.

Hydroelectric generation in PJM is composed of both run-of-river and pumped storage power plants. Run-of-river plants regulate between zero flow and full output, operating similarly to steam plants producing energy 24 hours a day. Pumped storage hydroelectric plants operate as pumping load off-peak to fill a reservoir, which releases water on-peak to generate electricity, with the resource modulating the water flow to provide regulation service. The hydroelectric model is very similar to the thermal model, with the exception of a 20 second time-constant in the low-pass filter, simulating a faster convergence to the desired output.

The energy storage response model is more complicated. An energy storage regulating resource, such as a battery or flywheel, cannot burn fuels to produce electricity, but instead charge and discharge energy from the bulk electric system itself through a power inverter. Energy storage devices can match a desired output within 200 milliseconds of receiving the regulating signal; therefore these resources do not need to model any dynamics, ramp limitations, or dead-bands. The limiting factor in operations is the charge storage capability itself, which is represented as the

number of minutes the battery takes to fully charge or discharge. First generation batteries typically run as 15-minute batteries, with newer devices capable of operating as 30-minute batteries or more. Assuming that a battery regulates around a 50 percent state of charge, the storage capability becomes:

$$0.5 \times \text{Total Regulation Capability (MW)} \times \frac{\text{Storage Size (Mins)}}{60 \text{ Minutes}}$$

When the accumulation exceeds the charge or discharge limits (100 percent or 0 percent of the storage capability), the output of the battery is set to 0 MW, until the device next receives a signal that discharges or charges the battery.

Due to similarities between energy storage and demand response, a separate model was not created for demand response resources. These resources respond very quickly to changes in the signal while also being limited by total energy as do energy storage resources.

In the scenario simulations, Regulation A response was composed entirely of thermal resources. However, Regulation D response was composed of 90 percent energy storage resources and 10 percent hydroelectric resources. These ratios were selected to match average operational behavior.

Historical balancing area data was available for calendar years 2015 and 2016, archived at two-second and 10-second intervals into an OSISoft PI Historian archive. This data contained historical regulation signals, the RTO ACE, and its components: actual and scheduled interchange, actual and scheduled frequency, the manual adder, the implemented frequency bias, L10, and epsilon values.

The challenge in analyzing the historical data was that the RTO ACE, interchange, and frequency all contained the effects of control from the historical regulation signals, so the team needed to generate an “uncorrected ACE” dataset that subtracted out the effects of regulating service. Regulating resources with various performance characteristics followed their respective regulation signals, which affects interchange, frequency, and ultimately RTO ACE. For the purposes of simplifying the analysis, frequency deviation was assumed to be constant between the historical and “uncorrected” models.

Interchange actual and schedule are determined by spot market interface pricing, and move as a function of the spread between prices on PJM's borders. Analysis of historical interchange did not show a significant correlation between flow and hour of day, day of month, or season of year. For the purposes of simplifying the simulation, interchange schedule was also assumed to be constant between the historical and “uncorrected” models.

The historical regulation signals were processed through the response model to generate historical response, which was then subtracted from the historical actual interchange to create an “uncorrected” interchange. New RTO ACE and control metrics were calculated as a baseline for improvement by the conditional neutrality regulation controller. Once the uncorrected RTO ACE was determined per historical day, the uncorrected ACE was fed through the conditionality neutrality controller with a closed-loop response calculation adjusting the next interval's actual interchange. In this way, each simulation's resultant ACE was unique to the control parameter selection. These unique resultant ACEs were then compared to determine the optimal parameters based on the control metrics.

With this data, several control metrics were calculated for benchmarking: NERC CPS-1, CPS-2, BAAL, and ACE-squared. For most comparisons and performance benchmarks, ACE-squared was used as the decision benchmark. This is because extreme ACE deviations are significantly more important than small ACE deviations and a controller which reduces extreme deviations at the cost of more, smaller deviations is materially better from an operations and ACE control perspective.

Seasonal Analysis for Regulation Requirement

As system conditions vary seasonally, the effects of this variation on regulation service were studied to establish how much regulation is appropriate to procure for the hours of the day in each season. The analysis took into account load ramping, CPS-1 scores and modeled regulation performance.

PJM's summer load curve is characterized by a single-peak sinusoid, with a continuous morning ramp peaking around 15:00, returning to a lower midnight valley load. Spring and fall shapes are characterized as transitional shapes, with a sharper morning ramp, flat afternoon load, and a sharp ramp to an evening lighting load peak, falling away to the overnight valley. The winter load shape is double-peaking, with sharp morning and evening ramps, low midday valleys and higher overnight valleys (driven by electric heating).




Figure 10: Seasonal Load Variation

The control metric performance was negatively correlated to periods of high load ramp (that is, delta load), not necessarily high load. PJM generates a five-minute interval load forecast spline used in the RTSCED application;

however this forecast is highly dependent on recent similar day behavior for its shape. Any inaccuracies in the short-term load forecast are covered by the regulation service as RTSCED moves to cover the difference in load, and higher rates of change in load expose the regulation service to higher volatility. It is expected that increasing the regulation requirement during periods of high ramp should improve overall control performance. This information was then used in order to determine the quantity of regulation required during various defined ramp and non-ramping periods. After simulations were run, results were filtered by specific ramp periods, and this information was used to determine what requirement levels would be necessary to ensure improved system control. Figure 11 below demonstrates the visible correlation between hours with low CPS-1 scores and high delta demand hours (high changes in load from one hour to the next).




Figure 11: Comparison of CPS-1 Scores and High Delta Demand

The modeled response seasonal simulations utilized data from the 2015 calendar year. The desire was to execute the controller against a variety of load patterns, interchange deviations and operating conditions to measure if the new controller could outperform the existing controller. The modeled results aligned with the defined ramp and non-ramping periods from the load and CPS-1 analysis performed. It was determined that a higher requirement would be needed for the ramping periods; however, less hours were needed to be classified as ramping in comparison to the on-peak requirement definition.

Regulation Market Requirement

The regulation requirement was split into two periods of the day since its inception as a cost-based service in 1998: defined as on-peak and off-peak periods. In 2012, the measured required MW was converted from assigned MW to benefits-factor and performance-adjusted effective MW, aligning with the PBR implementation. With the new requirement calculation allowing resource performance to factor into resource procurement, PJM was able to lower the regulation requirement, in comparison to historic values. Over the midnight “off-peak” period, from midnight to 5 a.m. local, PJM carried 525 MW of effective regulation requirement. In the remaining “on-peak” period from 5 a.m. local to midnight, PJM carried 700 MW of effective regulation requirement. A lower regulation requirement during the valley period is primarily viewed as a cost-savings action for PJM’s members, where the valley period is considered generally less prone to ACE disturbances than during the daylight period.

Beginning Jan. 9, 2017, the regulation requirement was updated to a seasonal, ramp/non-ramp definition. From the seasonal regulation requirement analysis, the results provided that CPS-1 performance would improve by carrying additional regulation during the periods of the day that exhibit the most change in load. In Table 1, Proposed Regulation Requirements, PJM identified hours of the day, per season, that exhibit higher rates of change in load, and proposed carrying additional regulation during those hours. In the remaining hours of the day, PJM reduced the amount of regulation requirement, in order to maintain roughly equivalent regulation MWh totals across the operating day. This defined regulation requirement will be evaluated quarterly by PJM staff for any necessary adjustments, based on new signal performance analysis.

Table 1 Proposed Regulation Requirements

Season	Months	Ramp-hours	Ramp MW	Non Ramp MW
Fall	September 1 – November 30	HE6 – HE8, HE18 – HE24	800	525
Winter	December 1 – February 29	HE5-HE9, HE17-HE24	800	525
Spring	March 1 – May 31	HE6 – HE8, HE18 – HE24	800	525
Summer	June 1 – August 31	HE6-HE14, HE19 – HE24	800	525

Scenario Generation

The controller tuning phase of the simulation involved executing a set of operating days over a variety of controller gains at historical total regulation capability values. The PID controller design effectively manages a high rate of convergence to a desired signal, but can exhibit overshoot and ringing if the controller is not tuned correctly. Ultimately, it was decided to remove the derivative gain to reduce the model to a PI controller; however, the parameter was retained in the production code.

In order to ensure that the controller was tuned properly, PJM simulated the controller in closed loop through thorough regression analysis. This process involved varying the various parameters of the controller simultaneously and studying the controller's results. After running the controller through rigorous testing, the following PID controller gains were determined:

Table 2 Conditional Neutrality Tuning Parameters

Controller Tuning Parameter	Value
Proportional Gain	1.0000
Integrator Gain	0.0100
Derivative Gain	0.0000
Regulation A Signal Time Constant	150
Regulation D Signal Time Constant	10

Once the conditional neutrality controller was considered stable, the total regulation capability values were varied to simulate the effects of carrying additional amounts of regulation capability. For a given operating day, the simulation tool iterated the total regulation capability for the Regulation A and Regulation D products, stepping from 0 MW to 1000 MW each in 50 MW increments, then limiting solutions between 400 MW to 1000 MW in total regulation capability. Additional steps were added for the first 120 MW of capability for each signal, in steps of 15 MW, to identify additional data points at the extremes. For each of these scenarios, the simulation tool calculated and retained average control metric values at hourly and daily aggregation levels.

With the new regulation signals in PJM's production environment, staff will continually monitor the signals and tune them in order to receive the desired ACE control out of the regulation product. The various tuning parameters mentioned above are available in PJM's EMS for fine tuning in order to ensure that the regulation signals closely align with expectations of resource responses and ACE control needs. The Regulation Market is very dynamic as technologies change and improve, and the ability of PJM staff to fine tune the regulation signals is essential to ensure that the signals match the needs of PJM for ACE control.

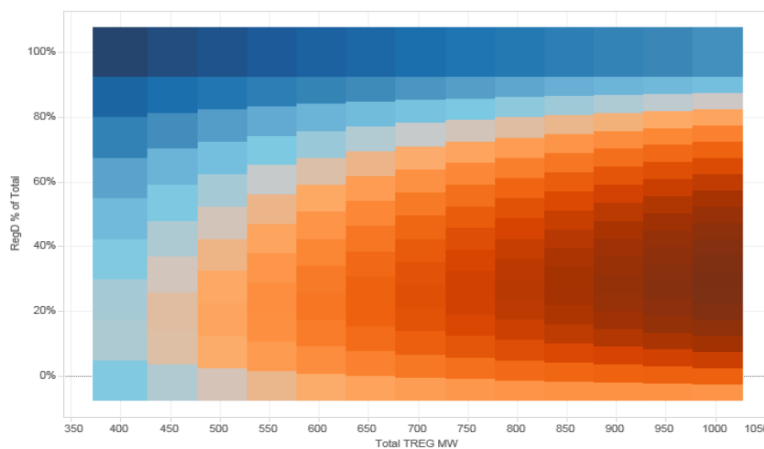


Figure 12: Example of Simulation Results Aggregated Across a Year

Rate of Technical Substitution Curve Formation

To begin defining the Rate of Technical Substitution (RTS) curve, “heat maps” were generated using control metric performance from the modeled response, at varying signal regulation capabilities, grouped by season and ramp/non-ramp types. These graphs demonstrated the net benefit to system control of increasing the amount of Regulation D that participated in the regulation service. To measure the trade-off between signal types, an isoquant was drawn to illustrate the equivalent performance across the varying capabilities, beginning at an intercept of (800 MW Regulation A, 0 MW Regulation D) for ramp periods, and (525 MW Regulation A, 0 MW Regulation D) for non-ramp periods.

PS Adj. MW	RegD																				
	0	50	100	150	200	250	300	350	400	450	500	550	600	650	700	750	800	850	900	950	1000
0									174,367	172,659	168,111	167,888	164,094	161,603	161,519	160,930	159,947	159,293	160,654	159,189	157,830
50								167,541	164,626	160,662	156,535	153,814	152,267	150,132	148,390	146,163	144,179	143,362	141,762	140,259	
100							152,987	149,345	143,896	139,212	136,218	135,452	133,066	130,925	129,160	126,764	124,555	123,703	122,968		
150					140,652	136,558	131,134	127,393	126,240	122,713	119,681	116,126	114,021	112,392	113,113	111,142	109,423				
200				135,332	130,335	125,500	119,326	117,428	114,040	111,217	107,032	103,539	102,589	100,143	99,094	98,235					
250			133,651	126,446	119,945	117,257	112,243	107,741	102,578	98,022	94,073	90,587	87,663	85,230	85,176						
300			131,801	122,671	114,920	111,583	106,004	101,124	94,878	90,055	85,831	82,183	80,084	79,216	77,392						
350		134,830	122,756	113,001	105,316	98,063	93,968	88,532	83,601	81,229	80,273	77,424	73,945	70,705							
400	142,229	125,505	112,860	103,198	96,828	92,933	86,882	81,680	79,098	74,819	70,806	67,126	63,742								
450	133,862	115,953	105,571	97,144	89,891	84,134	79,643	74,833	70,380	67,410	63,813	60,586									
500	126,260	110,686	100,514	90,475	83,407	76,947	71,747	67,115	63,032	59,342	56,075										
550	120,750	104,624	93,039	83,459	75,795	69,518	64,332	59,990	57,267	53,792											
600	113,706	98,083	86,872	77,597	70,314	64,363	59,407	55,236	51,759												
650	109,544	92,537	80,698	71,681	64,647	58,964	54,222	50,445													
700	103,605	87,027	75,473	66,715	59,953	54,531	50,189														
750	98,633	82,389	71,030	62,537	56,039	50,888															
800	94,506	78,555	67,411	59,070	52,726																
850	91,109	75,410	64,433	56,219																	
900	88,303	72,885	62,122																		
950	86,024	70,831																			
1000	84,215																				

Figure 13: Example of Heat Map for Varying Mixes of Regulation A and Regulation D

The isoquant results demonstrated that the system immediately benefitted by adding Regulation D to the regulation service. The new controller design slowed the Regulation A signal relative to the previous controller, so a pure Regulation A only service would have a slightly more difficult time controlling RTO ACE due to reduced desired ramping. Adding even small amounts of Regulation D results in a large reduction in total regulation that is required, which could result in a total cost savings to the Regulation Market.

At the other extreme, we found that system control would require large amounts of Regulation D capability to perform at equivalent capability due to the effects of charge storage limitations. Regulation A is able to maintain continuous full-raise and full-lower responses, due to their non-energy limited nature. Historically ACE deviations can sometimes last 30 to 40 minutes which requires there to be non-energy limited resources available to provide the regulation service. As a result, the battery storage size parameter significantly impacted the ability to perform to historical signals. Batteries with a 15-minute storage capability would perform well when RTO ACE was balanced, but would eventually drop out of regulation service if the RTO ACE deviation were to persist. The 30-minute batteries performed better, but were still not equivalent to traditional generators. The 45-minute and 60-minute batteries performed roughly equivalent, as the typical ACE deviation was shorter than the time it would take to deplete or over-charge a battery of this size.

Using the isoquant points of equivalent performance, a RTS curve was generated for use in the market clearing and pricing processes. Mathematically, this is the first derivative of the isoquant, which returns the rate by which we can replace Regulation A with Regulation D and still maintain consistent system control. The isoquant consists of the MW pairs of Regulation A and Regulation D which provide equivalent control.

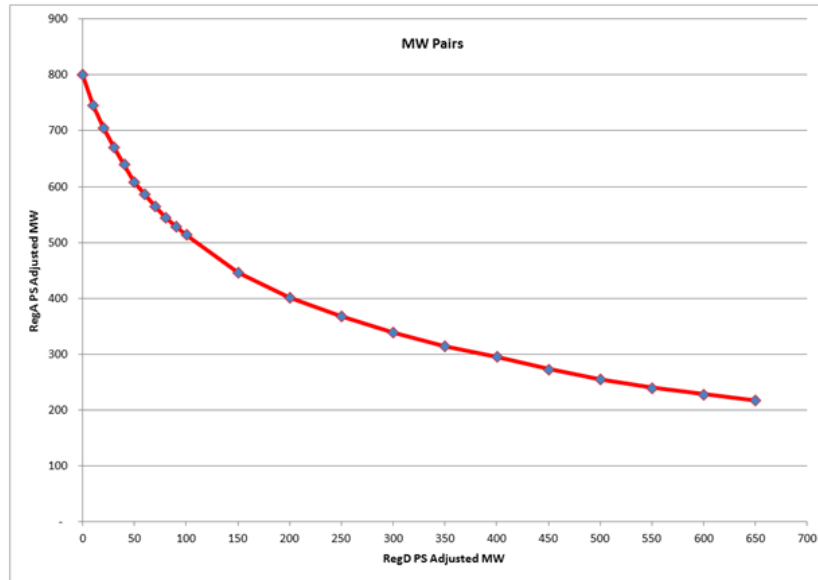


Figure 14: Isoquant - MW Pairs of Equivalent ACE Control

To best capture the rate of technical substitution at each point on the isoquant, the RTS curves were segmented and defined seasonally. The approach taken was to break each isoquant into four segments, then curve-fit a second order polynomial to the segment using a least-squared-error linear regression.




Figure 15: Isoquant Separated Into 4 Segments

The equations of the line fits for each segment were then run through a linear optimization to ensure that the ends of the polynomial segments converged with least error. With the parameters of the second order polynomial known, the first derivative of each line segment could be generated, which were then recorded as a four segment RTS curve. An example of this RTS curve is seen below in Figure 16.

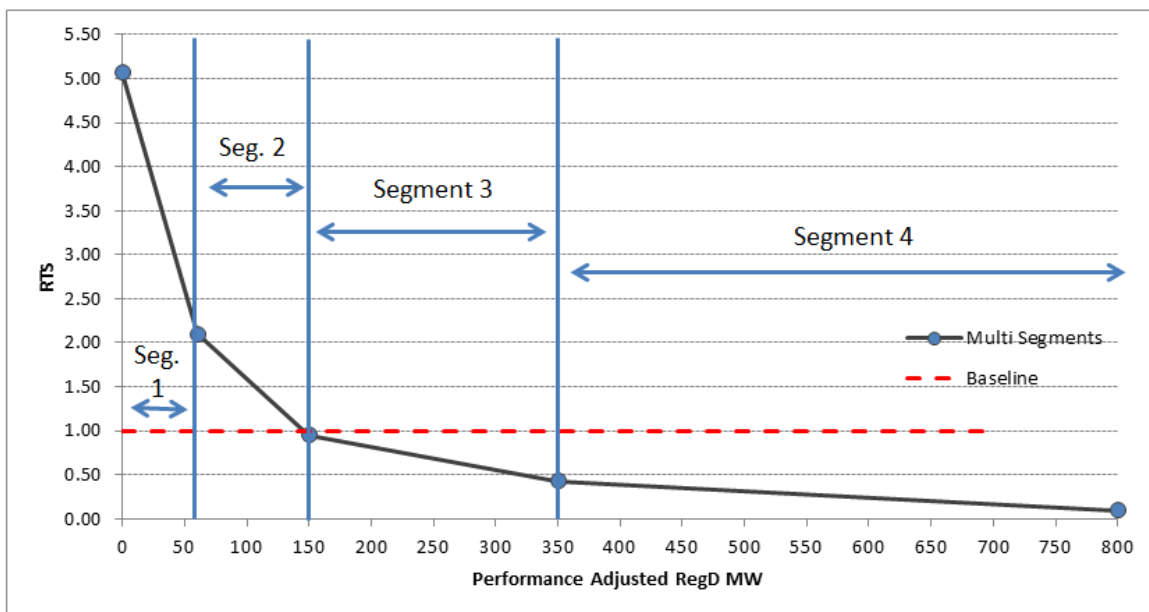


Figure 16: Segmented Rate of Technical Substitution Curve

Measurement and Verification of Substitution Curves

As PJM transitioned the new control signal to production, PJM has committed to updating the RTS curves to reflect real-world performance of regulation resources. The models used to simulate response to the signals are not as complex as actual plant response or human intervention. Thermal resources have internal processes that delay response to the control signals, or could respond better than expected if its ramp-rate exceeds its regulation offer. Energy storage resources can drop out of the regulation service, rebalance their state of charge values, and rejoin the service later in the same operating day. Additionally, the simulator modeled response as an aggregate of behavior, so individual resources would have more variation in performance scores, response profiles and storage capability.

As operating hours are accumulated and studied under the conditional neutrality controller, PJM will have the data required to know the actual performance to the regulation signals, as well as actual control metric values at varying mixes of Regulation A and Regulation D total capability ratios. In addition, PJM will continue to improve upon its resource response models in the simulation environment in order to accurately reflect real-world operational characteristics. Using operational data along with updated and improved study models, PJM will be able to generate control metric heat maps using actual values, and evaluate new isoquant performance lines. PJM will ensure that there are statistically significant numbers of samples at each total Regulation A and Regulation D capability point, so that staff will be able to re-execute the curve fitting methodology to produce new RTS curves for production that accurately represent real world performance.