California ISO
Shaping a Renewed Future

Minimum Effectiveness Threshold Report

March 1, 2010

Prepared by: Market Quality and Renewable Integration

California Independent System Operator

Minimum Effectiveness Threshold Report

Background

On March 23rd, the CAISO filed with the Federal Energy Regulatory Commission (FERC) the intention to use an effectiveness factor threshold setting of two percent in its market optimization software for determining which resources to utilize for redispatch to relieve congestion. Although some parties raised concerns about the effectiveness factor and the initial two percent setting, the CAISO explained how it had arrived at the two percent value after trying and assessing the impacts of alternative values in market simulation, and had determined that two percent was a reasonable starting value that would closely reflect the actual operating practices of grid operators and would balance the potential for undesirable impacts of either too high or too low a value. At the same time the ISO committed to monitor and assess the impacts of the two percent value in market production to be able to determine, based on actual market experience, whether the value of the threshold should be changed. On June 10, 2009, FERC accepted the 2% threshold as a reasonable starting value and directed the CAISO to conduct analysis and report its findings by March 1, 2010. The ISO provides the present report in compliance with FERC's direction.

Stakeholders expressed two areas of concern with regard to the effectiveness threshold. The first concern was whether the effectiveness threshold would result in any systematic operational or market anomalies or performance issues. In particular, whether the two percent value would either limit the ability of the software to effectively manage congestion or fail to prevent the software from utilizing large MW adjustments of ineffective resources. The two percent threshold was selected to reflect prudent and accepted operational practice in the sense that operators would redispatch resources to relieve congestion on a constraint only if those resources were relatively close to the constraint so that changes in their dispatch levels would be effective in modifying the flow on the constraint. In other words, the operators would avoid managing congestion using large redispatch adjustments on resources distant from and ineffective on the congested constraint in order to minimize the total or gross amount of redispatch. Setting too low a value for the effectiveness threshold could result in large redispatch movements with little or no congestion relief benefits, which in turn could negatively impact the ability of the operators to maintain system balance or control. On the other hand, too high a setting of the effectiveness threshold could hamper congestion management by excluding operationally sound potential redispatch adjustments and forcing the optimization software to relax a constraint that could have been relieved. In this report the ISO provides an analysis based on operational experience of the actual impact of the two percent effectiveness threshold.

The second concern expressed by stakeholders was that the threshold may cause inconsistency of pricing in the day-ahead market between an Aggregate Pricing Node (APNode) used to price and settle demand at a Default Load Aggregation Point (DLAP) as compared to the optimal dispatch level of the corresponding aggregated resource (Anode) demand bid submitted at the DLAP. The APNode price is the weighted average price of all constituent Pricing Nodes (PNodes) weighted by the quantity of load at each PNode. Hence its value can be affected by any redispatch adjustments the software makes to resources at individual PNodes that are effective in relieving congestion. In contrast, the schedule for an aggregated DLAP resource bid into the day-ahead market is determined based on the effectiveness of adjustments to the aggregated DLAP resource relative to the congested constraint. As a result, depending on the location of congestion and effectiveness of the DLAP resource, the DLAP resource may not be above the effectiveness threshold and therefore would not be adjusted to relieve congestion, while resources at some of its constituent PNodes could be adjusted to relieve congestion on the same binding constraint. In such an instance the ANode price associated with the market schedule of the aggregated DLAP resource would reflect the fact that the resource was not adjusted for congestion relief, while the APNode price as the weighted average of constituent PNode prices and schedules would reflect the fact that adjustments were made at some of those PNodes for relieving congestion on the same constraint. Hence the two prices could be inconsistent and, of specific concern to the stakeholders who raised this issue, the APNode settlement price may differ from the ANode price used to clear the day-ahead schedule of the aggregated DLAP resource. Below are the results of the ISO's analysis to assess the frequency and impact of this potential inconsistency.

Effectiveness Threshold

The measure of how effective a resource at a particular location is to change the flow on a constraint is referred to as the shift-factor or the power transfer distribution factor. The shift factor at a node for a given flowgate represents the change in the amount of MW flow through the flowgate in the reference direction given an assumed additional MW injection at the node and a corresponding withdrawal from the reference slack. The shift factor of a node for a given flowgate is between -1 and +1 and depends on the location of the node relative to the flowgate within the network, the choice of the slack, and the reference direction of the flowgate. For two given nodes A and B at different locations, the difference between their shift factors (shift factor A minus shift factor B) is the change in flow through the flowgate in the reference direction with an assumed 1MW additional injection into bus A with the corresponding withdrawal from bus B. Though the shift factor of a bus for a given flowgate is different under different choice of slack, the difference in shift factors between two buses is the same irrespective of the choice of the slack. Thus the difference between the shift factors of nodes A and B for a given flowgate measures the effectiveness of relieving congestion on the flowgate through an increase in the dispatch at one node and an equal decrease in the dispatch at the other node. Positive difference means that additional flow is along the flowgate reference direction for injection at A and withdrawal at B.

Assuming a system dispatch that would result in a flow limit violation of a flowgate in the reference direction that requires a re-dispatch to resolve the violation, reducing 1MW from generator at node A with shift factor SFA and increasing 1MW for generator at node B with factor SFB, where SFA > SFB, the flow on the flowgate is reduced by the amount (SFA – SFB) MW while maintaining power balance. Thus, the difference in shift factors at the two locations reflects the combined effectiveness of the incremental and decremental adjustments, representing the amount of MW flow reduction for 1 decremental MW of generator at bus A and 1 incremental MW of generator at bus B.

The effectiveness threshold is designed so that resources at a pair of nodes will be adjusted to relieve a flow violation on a particular flowgate only if they both have effectiveness factors greater than the threshold for that flowgate. Since the start-up of the new markets the ISO has adopted the effectiveness threshold value 0.02 or 2%. As such, for any given flowgate, the market software sets the shift factor to zero for a resource that is less than 0.02 with respect to the distributed load slack.

In essence, the set of resources originally having shift factors in the range (-.02, +.02) for this flowgate will be excluded from consideration by the software for re-dispatch adjustment to resolve a flow violation on the flowgate.

The most effective adjustment to correct a flow limit violation for a given flowgate is using the pair the generators that have the largest difference in their shift factors. In fact, this is the pair that requires the least amount of MW adjustment per MW of flowgate flow reduction. On the other hand, for a pair of generators whose difference in effectiveness is small, large MW adjustments are needed per MW flow reduction, hence these are ineffective. One limitation of the entire effectiveness factor approach is that, because it is applied to the shift factor of each node individually, it cannot prevent the software from making adjustments to a pair of resources whose difference in effectiveness – hence their combined effectiveness on the flowgate – is very small even though their individual effectiveness values both exceed the threshold. As discussed further below, when this phenomenon occurs the problem is not with the level of the effectiveness threshold but rather with a basic principle of how the effectiveness threshold works.

Evaluation of Operational Experience

Below is an evaluation of actual operational experience since the start of the new market design that illustrates the extent to which the application of the effectiveness threshold either: 1) reduced the ISO's ability to manage and resolve congestion, or 2) was not effective in limiting ineffective congestion relief.

East Nicolas to Rio Oso 115 kV and Palermo to Honcut Junction 115 kV Constraints

The ISO's review of market results to date shows that there have been instances in which the effectiveness threshold has excluded adjustments to resources that could have been effective in providing congestion relief even though their effectiveness factors were low. Although these instances have been limited, their occurrence indicates that a closer examination is warranted, as is presented below.

One instance involves constraints in the Sacramento Valley that were discussed in the ISO's first quarterly Post-Implementation Report prepared by the ISO's Department of Market Services, which was filed with FERC on July 30, 2009. The East Nicolas to Rio Oso 115 kV line was relaxed by 0.2 MW in the day-ahead market on May 2, 2009, producing a pricing run shadow price of $4740.65/MWh. It was again relaxed by 0.8 MW on May 14, 2009, producing a pricing run shadow price of $500/MWh. Also on May 2, 2009, the nearby Palermo to Honcut Junction 115 kV line had a maximum shadow price of $725.71/MWh, and was enforced without relaxation. These constraints resulted from outages in the higher voltage transmission system running north-to-south through the Sacramento Valley; the ISO had multiple days around this time when this 115 kV transmission system had significant congestion costs due to the north-to-south flows, until the ISO was able to later identify a remedy of transmission circuit switching to relieve this congestion.

Further examination shows that the largest contributor to north-to-south flows through the Sacramento Valley, which is imports from the northwest across the Pacific AC Intertie (PACI), has a power transfer distribution factor (equivalent to effectiveness factor) ranging from just above 2% to just below 2%. In the hours of the greatest congestion impact on these constraints, PACI imports were less than 2% effective for managing these constraints, which led them to be excluded from participation in congestion management. After replicating the conditions reported in the quarterly Post-Implementation report using the current version of the ISO's market software, re-running the May 2 case with an effectiveness factor of 1% shows that the East Nicolas to Rio Oso constraint could have been managed without constraint relaxation, reducing the highest LMPs by about 7%, and reducing the ISO's overall congestion costs in the day-ahead scheduling run's optimization by more than 30%.

Figure 1:

Frequency Distribution of Resource Shift Factors for
32212_E.NICOLS_115_32214_RIO OSO _115_1_1

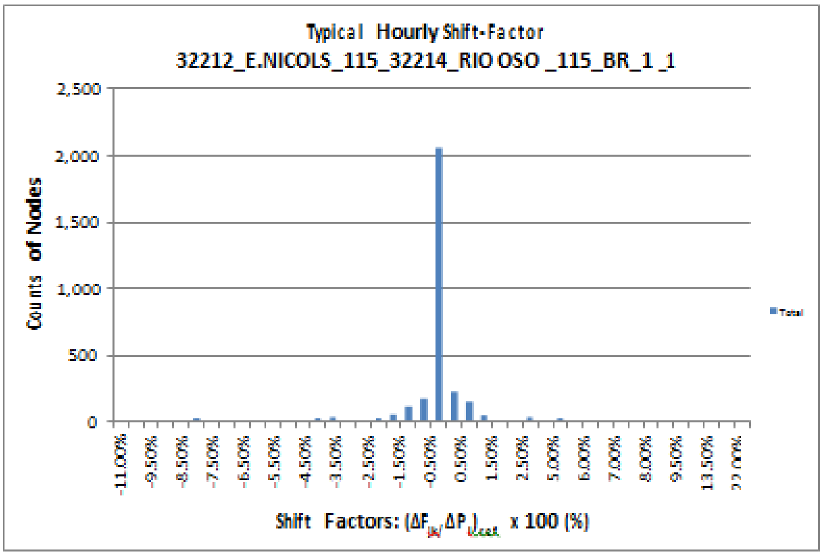

VICTVL_BG Constraint

One example of a situation in which the effectiveness threshold may have limited the ISO's ability to resolve congestion is the VICTVL_BG constraint. This VICTVL_BG constraint is at the boundary of the ISO and its neighboring LADWP Balancing Authority Area, and one result of the network topology in the area is that even resources electrically close to the constraint have very low shift-factors relative to the distributed load reference bus. As a result, resources that operators would have expected to be effective and were historically used for flow relief of VICTVL_BG based on operating procedures, instead have low shift factors. As shown in Figure 2, many of the shift factors for resources historically used for managing the VICTVL_BG constraint are at or below the 2% effectiveness threshold. Any resource whose shift-factor is below the effectiveness threshold will have its shift-factor replaced with a zero shift-factor relative to the constraint. As a result, such resources may be adjusted to counter movement on effective resources to maintain power balance but such resource movements may not be prioritized nor priced relative other resources that also have shift-factor below the effectiveness threshold but may be relatively less effective. The typical shift-factors for the VICTVL_BG provided in Figure 2 also illustrates that increasing the effectiveness threshold could impact significant amount of resources that are very close to the 2% threshold. While Figure 2 reflects a typical hour's set of shift factor, depending on the makeup of the distributed reference or transmission outages on the system, the shift-factor may change such that a resource may be effective one hour but may become ineffective another hour.

Figure 2:

Frequency Distribution of Resource Shift Factors for VICTVL_BG

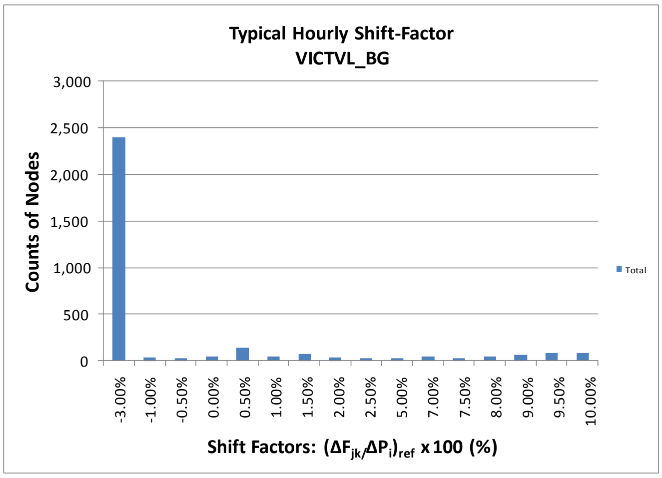

HUMBOLDT_BG Constraint

Another example where the application of a higher effectiveness threshold could affect the pool of generating resources available to relieve the congestion is the HUMBOLDT_BG constraint. As shown in Figure 3, there are there are few resources that are highly effective with -.995 (99.5% effective) shift-factors and that if increased could reduce the congestion relative to the distributed reference. Any increase in such a resource would, however, require offsetting decreases in other resources to maintain system balance, and as Figure 3 shows there are large number of generating resources in the rest of the system outside the Humboldt area that have shift-factors less than the effectiveness threshold relative to the distributed load reference and would be excluded from participating in relieving the constraint. Resources outside the Humboldt area that have shift-factors that fall below the threshold will have their shift-factors replace by zero. As a result, such resources may be adjusted relative to the effective resources in the Humboldt area to maintain a power balance but such movement will not reflect the relative effectiveness[1] of such resources in resolving the constraint that fall below the threshold and nor will such adjustment be reflected in the congestion component of the resources. In the Humboldt case the fact that there are a few resources that are effective while the balance of the resources in the system are not effective has little impact on the ISO's ability to resolve the constraint, however it does result in prices on resources that are below the effectiveness threshold to not reflect a resource's impact on relieving congestion.

[1] Differences in losses may result in differentiation of resource adjustment for resources that have been deemed ineffective.

Figure 3:

Frequency Distribution of Resource Shift-Factors for HUMBOLDT_BG

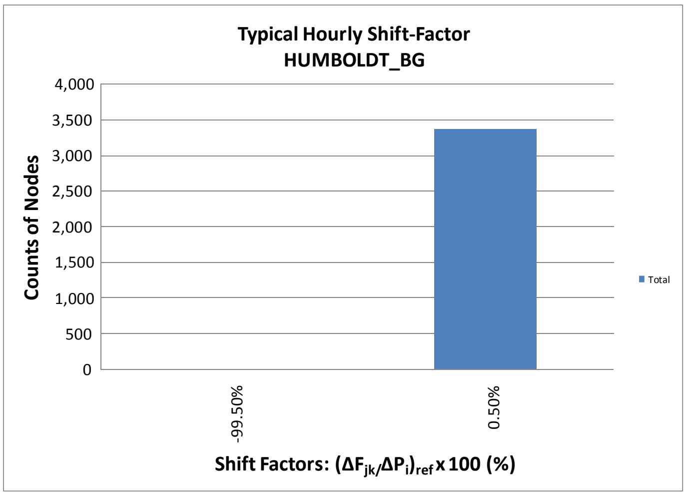

SDGE_CFEIMP_BG or SCE_PCT_IMP_BG Constraint

The SDGE_CFEIMP_BG (Figure 4) or SCE_PCT_IMP_BG (Figure 5) constraints provide examples where application of the effectiveness threshold to individual resources does not prevent the software from potentially performing ineffective re-dispatch. In the case of these constraints the shift factors that are generally above the effectiveness threshold that can be grouped in to one group off effective resources with similar negative shift factors on one side of the constraint and another group with positive shift factors on the other side of constraint. In these cases, if there is sufficient resource capacity to re-dispatch on each side of the constraint using effective shift-factors that are above the effectiveness threshold, the software solutions are generally stable and effective. Ineffective re-dispatch combinations can occur, however, when dispatch capability (i.e., ramping or maximum generating capacity) is exhausted on one side of the constraint and the only remaining capacity available is from resources with shift factors that may individually be well above the threshold but are only slightly different from each other. In these situations the combined effectiveness of the adjusted resources can be quite small, thus requiring large MW amounts of re-dispatch.[2] This situation is rarely observed in the day-ahead market because of the greater flexibility of the software to commit additional resources and the lower likelihood that ramping limitations will become binding. However, this has occurred in the Real-Time Market under some conditions.

---

[2] This is similar to the phenomenon observed related the use of lossless shift factors documented in Technical Bulletin: http://www.caiso.com/Documents/TechnicalBulletin-Comparison-LossyversusLoselessShiftFactorsinISOMarketOptimizations.pdf

Figure 4:

Frequency Distribution of Resource Shift-Factors for SCE_PCT_IMP_BG

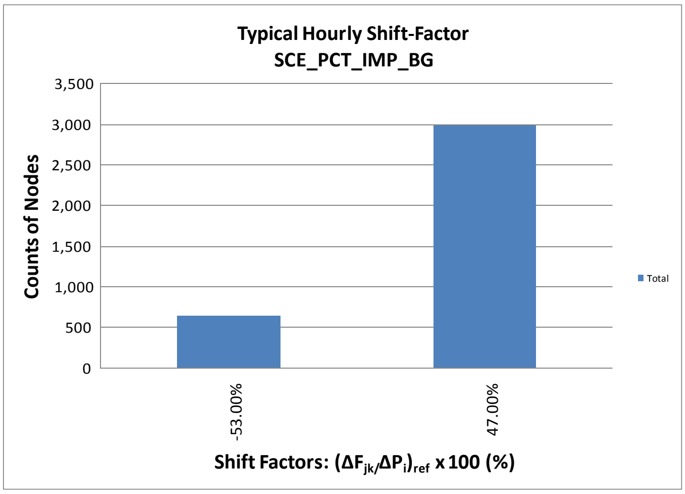

Figure 5:

Frequency Distribution of Resource Shift-Factors for SDGE_CFEIMP_BG

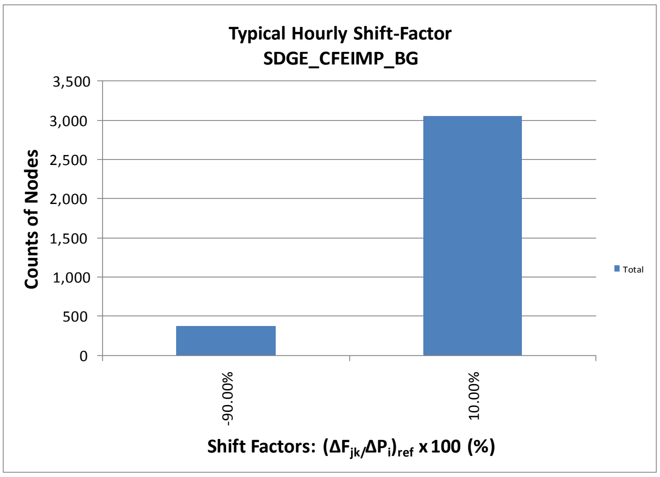

<u>APNode versus ANode Pricing Analysis</u>

An APNode price is determined based on the weighted average price of all constituent Pricing Nodes (PNodes) weighted by the quantity of load in each PNode. Hence its value can be affected by any redispatch adjustments the software makes to resources at individual PNodes that are effective in relieving congestion. In contrast, the schedule for an aggregated DLAP resource (also referred to as an ANode) bid into the day-ahead market is determined based on the effectiveness of adjustments to the aggregated DLAP resource relative to the congested constraint.[3] As a result, depending on the location of congestion and effectiveness of the DLAP resource, the DLAP resource may not be above the effectiveness threshold and therefore would not be adjusted to relieve congestion, while resources at some of its constituent PNodes could be adjusted to relieve congestion on the same binding constraint. In such an instance the ANode price associated with the market schedule of the aggregated DLAP resource would reflect the fact that the resource was not adjusted for congestion relief, while the APNode price as the weighted average of constituent PNode prices and schedules would reflect the fact that adjustments were made at some of those PNodes for relieving congestion on the same constraint. Hence the two prices could be inconsistent and, of specific concern to the stakeholders who raised this issue, the APNode settlement price may differ from the ANode price used to clear the day-ahead schedule of the aggregated DLAP resource. Figure 6 illustrates the scenario where the APNode price is higher such that it intersects with the DLAP bid curve at a lower MW level than the level that the ANode cleared at. As the shift-factor effectiveness threshold were reduced the ANode and APNode clearing MW and price would converge.

Figure 6

Anode vs APNode Clearing Level on DLAP Demand Bid Curve
where APNode Price is Inconsistent with Anode Cleared level

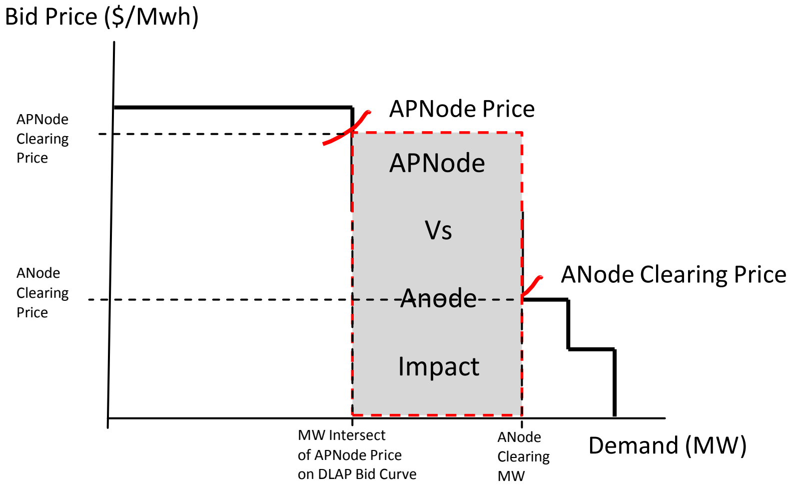

Below is an analysis to assess the frequency and impact of inconsistencies between these prices. Table 1 provides a comparative analysis of the frequency,

[3] It is important to note that the potential inconsistency discussed here can occur only in the day-ahead integrated forward market, because that is the only market that accepts and can clear demand bids that utilize DLAPs.

magnitude, and impact of APNode and Anode price differences. Figures 7, 8, and 9 provide an overview of the frequency and magnitude of the APNode vs. Anode differences. In each of these figures the left-hand graph shows the entire range of price differences, while the right-hand graph shows a magnified view of the five percent of hours in which the price differences were the greatest.

Table 1:

Summary of Comparison of APnode vs. Anode

(April 1, 2009-Febuary 20, 2010)

Comparison	PGAE_DLAP	SCE_DLAP	SDGE_DLAP
% of Day-Ahead hours with ABS(APNode - Anode) < $.01	86.8%	94.1%	98.0%
% of Day-Ahead hours with ABS(APNode - Anode) < $1.0	98.78%	99.8%	100%
Maximum Positive Price Difference (APNode>Anode)	$23.28	$45.15	$0.14
Maximum Negative Price Difference (APNode<Anode)	-$1.68	-$0.29	-$0.15
Estimated settlement impact of APNode and Anode price difference[5]	$75,068	$228,225	$0

Figure 7:

PGAE_DLAP APNode / Anode Difference Frequency

a. All Hours b. Top 5% of Hours

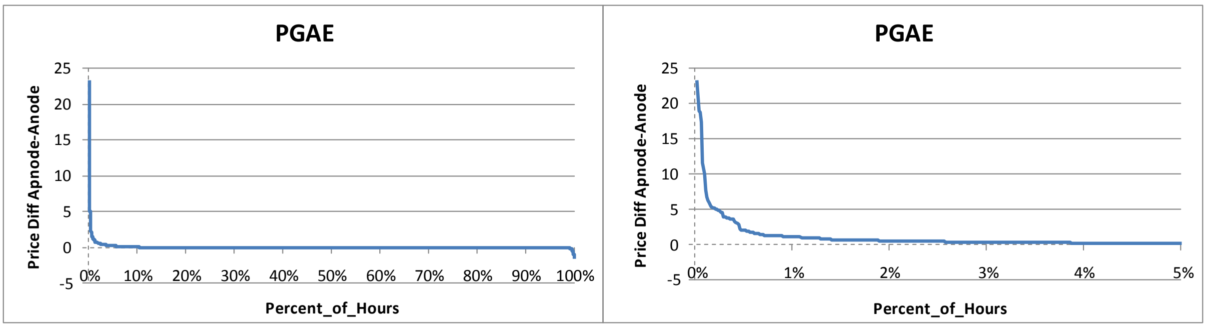

Figure 8:

[5] Estimated settlement impact was determined based the APNode price multipled by the MW level difference between DLAP cleared in the market and the maximum level that would have on the DLAP bid curve intersecting with the APNode price.

SCE_DLAP APNode / Anode Difference Frequency

a. All Hours

b. Top 5% of Hours

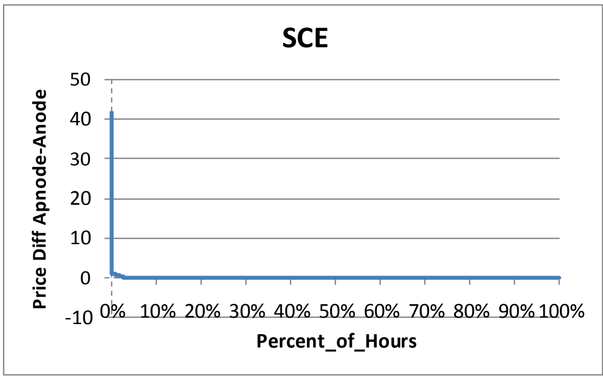

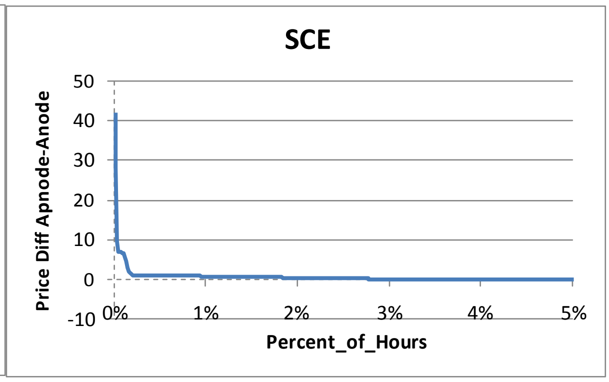

Figure 9:

SDG&E_DLAP APNode / Anode Difference Frequency

a. All Hours

b. Top 5% of Hours

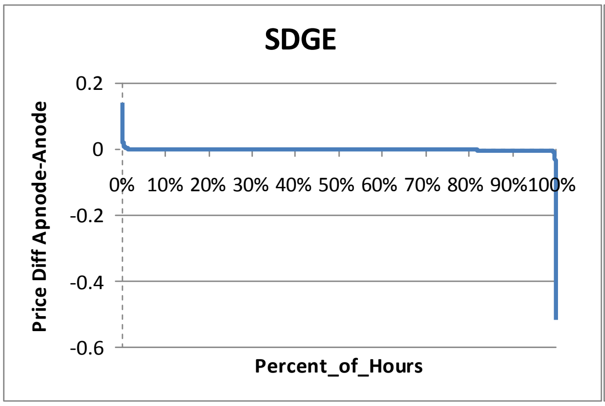

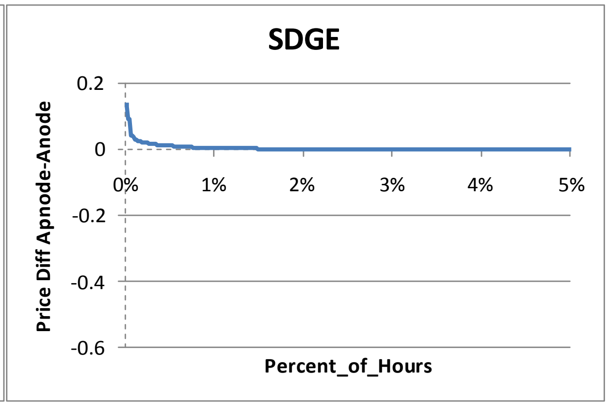

Conclusion

To date the ISO has not found sufficient evidence to indicate a need to either increase or decrease the two percent effectiveness threshold, at this time, with which the new markets have operated since the April 1, 2009 start-up. In reaching this conclusion the ISO has carefully monitored and analyzed the data related to the two concerns raised by parties in their previous comments on the two percent threshold value, namely, the potential to distort congestion management by utilizing ineffective adjustments or precluding potentially effective adjustments, and the potential to cause discrepancies between APNode and ANode DLAP energy prices in the day-ahead market. The ISO's experience does reveal, however, that further observation and consideration are needed regarding a limitation in the effectiveness threshold approach itself, which allows the market software to adjust multiple resources that are highly ineffective in combination even though the individual effectiveness of each resource is above the threshold. The anecdotal evidence reveals that in some cases, the single resource effectiveness threshold may reduce the ability of the market software to effectively manage and price congestion. Furthermore, applying the effectiveness threshold to individual resources may not always be successful in limiting highly ineffective dispatch outcomes. As noted earlier, however, these concerns arise not from the specific level of the effectiveness threshold but from the approach of applying the threshold to individual resources independently of one another. To date after much consideration the ISO has not been able to identify a workable approach that would address this phenomenon. Lastly, while the use of 2% effectiveness threshold has at times resulted in differences between the APNode and Anode DLAP prices, the overall impact of such differences has been extremely limited.