Final Report
Analysis Track Testing of CAISO MRTU Pricing and Dispatch

Scott M. Harvey, Matthew Kunkle, Benjamin Hagberg and Shaun Glassman[1]

October 20, 2008

EXECUTIVE SUMMARY

LECG was retained by the California ISO to review the results of the CAISO's analysis track testing of its MRTU dispatch and pricing software for the day-ahead market, real-time pre-dispatch, and real-time dispatch. The purpose of the CAISO analysis track testing was to test the software that has been developed for operating the CAISO MRTU electricity markets under the pricing rules described in the CAISO tariff.

The analysis track testing allowed the CAISO and LECG to test the day-ahead market (IFM), real-time pre-dispatch and unit commitment (RTUC) and real-time dispatch (RTD) software and associated pricing modules that will be used to coordinate the CAISO day-ahead and real-time electricity markets under MRTU. LECG's assignment was to assess the economic consistency of the results produced by the day-ahead market, real-time pre-dispatch and real-time dispatch software and pricing modules, and verify that any inconsistencies identified did not arise from errors in the calculation of settlement prices or reflect substantial deviations from the least-cost dispatch. This report summarizes the results of our analysis to this point in the testing process.

The CAISO MRTU software systems are designed to make optimal commitment and dispatch decisions based on market participant bids, subject to a variety of physical equipment constraints and power system reliability considerations. We have taken advantage of this property in testing the software, as the prices and schedules developed by the software should satisfy a series of internal consistency and equilibrium conditions if the software is functioning properly in both calculating prices and developing schedules. Inconsistencies are then reviewed to determine whether they reflect an error in the calculation of prices, an error in the process of determining the least cost unit commitment and dispatch, a data export error or other factors. In addition to the analysis track cases, this report also covers the results of our review of a series of cases used by the CAISO to test whether particular features of the software operated correctly, or to test

[1] Scott Harvey (sharvey@lecg.com) is a director with LECG. Matthew Kunkle is a managing consultant, Benjamin Hagberg an associate and Shaun Glassman a research analyst. Alexis Maharam and Christine Offerman were involved in the preparation of the initial and interim reports but did not participate in the preparation of the final report.

whether the software operated correctly during certain kinds of conditions. These are referred to as Class B cases, and are discussed in Section III of this report.

Since the completion of our interim report issued on July 1, 2008, we continued our evaluation, focusing on issues with the implementation of forbidden regions, intertie price calculations, energy limits, and the cascading of ancillary service prices, and testing cases that included binding nomogram and contingency constraints, wheeling transactions, and multi-hour block transactions. All of these issues are resolved in the most recent test cases. We carried out a complete review of these recent test cases to ensure that no new problems arose or old problems returned as a result of the software changes. Based on the analyses we have performed, we have not observed substantial unresolved problems that would prevent the CAISO software systems from calculating prices consistent with the CAISO tariff and LMP pricing methodology and have not observed material unresolved problems that would prevent the software systems from committing and dispatching load and generation based on least bid cost. Our review of the class B cases found that the features of the CAISO software being tested in these cases performed as intended in each instance.

Final Report
Analysis Track Testing of CAISO MRTU Pricing and Dispatch

Scott M. Harvey, Matthew Kunkle, Benjamin Hagberg and Shaun Glassman[2]

October 20, 2008

I. INTRODUCTION

LECG was retained by the California ISO to review the results of the CAISO's analysis track testing of its MRTU dispatch and pricing software for the day-ahead market, real-time pre-dispatch, and real-time dispatch. The purpose of the CAISO analysis track testing was to test the software that has been developed for operating the CAISO MRTU electricity markets under the pricing rules described in the CAISO tariff.

The analysis track testing allowed the CAISO and LECG to test the day-ahead market (IFM), real-time pre-dispatch (RTUC) and real-time dispatch (RTD) software and associated pricing modules that will be used to coordinate the CAISO day-ahead and real-time electricity markets under MRTU. LECG's assignment was to assess the economic consistency of the results produced by the day-ahead market, real-time pre-dispatch and real-time dispatch software and pricing modules based on the analysis track test cases provided to us by the CAISO and verify that any inconsistencies identified did not arise from errors in the calculation of settlement prices or reflect substantial deviations from the least-cost dispatch. This report summarizes the results of our analysis to this point in the testing process. This report describes the scope of the testing we have carried out and that is covered by this report, relative to other testing carried out by the CAISO that may be reported elsewhere.

The CAISO MRTU software systems are designed to make optimal commitment and dispatch decisions based on market participant bids, subject to physical equipment constraints and power system reliability considerations. We have taken advantage of this property in testing the software, as the prices and schedules developed by the software should satisfy a series of internal consistency and equilibrium conditions if the software is functioning properly in both calculating prices and developing schedules. Inconsistencies are then reviewed to determine whether they reflect an error in the calculation of prices, an error in the process of determining the least cost unit

[2] Scott Harvey (sharvey@lecg.com) is a director with LECG. Matthew Kunkle is a managing consultant, Benjamin Hagberg an associate and Shaun Glassman a research analyst. Alexis Maharam and Christine Offerman were involved in the preparation of the initial and interim reports but did not participate in the preparation of the final report.

commitment and dispatch, a data export error or other factors. These tests do not assess the accuracy of the underlying transmission grid model, the accuracy of the powerflow solution, or the accuracy of the generation shift factors or loss penalty factors used in the dispatch. This test process also serves to identify gaps in the data saved by the MRTU software that would hinder price validation and analysis of market performance under MRTU operation. In addition to the analysis track cases, this report also covers the results of our review of a series of cases used by the CAISO to test whether particular features of the software operated correctly, or to test whether the software operated correctly during certain kinds of conditions. These are referred to as Class B cases, and are discussed in Section III of this report.

Since the completion of our interim report issued on July 1, 2008, we continued our evaluation, focusing on issues with the implementation of forbidden regions, intertie price calculations, energy limits, and the cascading of ancillary service prices, and testing cases that included binding nomogram and contingency constraints, wheeling transactions, and multi-hour block transactions. All of these issues are in the most recent test cases. We carried out a complete review of these recent test cases to ensure that no new problems arose or old problems returned as a result of the software changes. Based on the analyses we have performed, we have not observed substantial unresolved problems that would prevent the CAISO software systems from calculating prices consistent with the CAISO tariff and LMP pricing methodology and have not observed material unresolved problems that would prevent the software systems from committing and dispatching load and generation based on least bid cost. Our review of the class B cases found that the features of the CAISO software being tested in these cases performed as intended in each instance.

II. TESTS OF MRTU SOFTWARE

We conducted several rounds of tests on CAISO analysis track test cases[3] in order to validate the price calculations, unit commitment and dispatch solutions and identify potential economic inconsistencies in the prices and schedules produced by the CAISO day-ahead market, and real-time pre-dispatch and unit commitment software, and the real-time dispatch software (RTD).

The initial round of analysis track testing entailed evaluation of 51 cases. The results of our evaluation were reviewed with the CAISO and problems with the initial software version were identified for correction by Siemens. A second partial round of analysis track testing was carried out between November 20 and December 18 on 7 IFM and 5 RTUC cases to assess whether the problems identified in the initial round of testing

---

[3] A description of Analysis Track Testing scope of work can be found on the CAISO website at: http://www.caiso.com/1f8d/1f8d80dc2c580.pdf

had been corrected and to verify that the process of correcting these problems had not introduced new problems. During January, February, March and early April a single IFM and RTUC base case were tested and retested on each software patch until all anomalies were resolved. Finally, a third round of analysis track testing was carried out beginning on April 2, 2008.[4] The April-June 2008 testing covered 11 IFM cases, 11 RTUC cases, and 4 RTID cases, including all of the situations tested by the original 51 cases and some additional situations.[5] In addition, during June 2008 3 IFM and 3 RTUC cases were rerun using the then most recent software version and re-evaluated to verify that previously identified errors have been corrected. Finally, during July, August and September we tested and retested software functionality relating to forbidden regions wheeling transactions, export pricing, block transactions, energy limits, ancillary service pricing, and binding contingency and nomogram constraints.

A. Test Methodology

We carried out a series of tests on the analysis track cases in order to validate the price calculations, unit commitment and dispatch solutions and identify potential economic inconsistencies in the prices and schedules produced by the MRTU software. We then reviewed the inconsistencies to determine whether they reflected errors in the calculation of prices or material departures from the least cost unit commitment and dispatch. As described below, our tests of the day-ahead market and predispatch software included a review of energy and ancillary service prices and schedules. We also tested the energy schedules produced by the real-time dispatch software. The ancillary services included in the testing were regulation up, regulation down, 10-minute spinning reserves and 10-minute (or non-spinning) reserves.

Our tests of the economic consistency of the day-ahead market and real-time predispatch software verified that the correct relationship existed between resource and load schedules (including imports and exports), energy and ancillary service prices, and energy and ancillary service bid and offer prices.[6] Our tests of the economic consistency of the real-time interval dispatch software verified that the correct relationship existed between the schedules of dispatchable resources, their offer prices, and real-time energy

---

4 The 26 cases tested in the third round of analysis track testing were designed to cover all 51 of the cases included in the initial round of testing. The initial 51 cases included many cases with very minor differences to facilitate identification of the source of software performance issues. This level of disaggregation was not necessary in the third round which was intended to verify that the software was operating as intended.

5 Some cases have been reviewed at more than one MIP gap level and run on more than one computer system.

6 These tests were generally based on the prices and resource schedules in the pricing dispatch, rather than the physical dispatch. However, in the case of resources not eligible to set prices, we verified that their schedule in the pricing dispatch was identical to their schedule in the physical dispatch and that their schedule in the physical dispatch was consistent with the prices in the physical dispatch.

prices. In those instances in which the tests detected anomalies, we worked with CAISO staff and/or Siemens to determine whether the anomalies reflected incorrect software performance, incorrect data, reflected false positives (i.e., the anomalies in fact reflected the correct operation of the software), or reflected isolated instances of imperfect optimization over integer variables (termed MIP gap, discussed further below).[7] As a general matter, the tests were designed to find flaws in the dispatch and pricing logic, not to systematically test the underlying transmission system data or powerflow solutions. In practice, however, the tests also identified other types of flaws, such as incorrect data modeling for some generation and connectivity to the network model. These data issues were either addressed during testing or will be corrected as the model and master file are updated.

1. Replicate LMP Energy Prices

For the day-ahead market, real-time predispatch and real-time interval dispatch software, this test verifies that we can replicate the calculation of the LMP prices from the components. For each generator and load bus for which a LMP price is calculated (Pnodes) and each aggregated generator or load bus for which a LMP price is calculated (Apnodes) we verify that we can calculate the LMP price from the distributed reference bus price, constraint shadow prices and penalty factor and generation shift factors for that location, as calculated by the day-ahead market, real-time predispatch and real-time dispatch models based on the underlying transmission grid model. This test not only serves to identify potential errors in the calculated prices, it also serves to verify that the correct shift factors, constraint shadow price, and penalty factors are being saved and exported, and thus will be available for analysis of performance issues.

2. Validate LMP Energy Prices Based on Marginal Generators

This test verifies that the appropriate number of marginal generators can be identified for the determination of LMP energy prices for each hour (in the day-ahead market), each 15-minute interval (in the real-time predispatch) and each 5-minute dispatch interval in the real-time dispatch. In an interval with no binding transmission constraints, at least one generator or price capped load bid should be on the margin (i.e., partially dispatched and not ramp constrained or at an upper or lower operating limit)[8] and the LMP price at

---

7 False positives in the initial testing were in part an intentional design feature to facilitate identification of problems and in part reflected limitations in the software output which in some cases made it difficult to verify that resource schedules were affected by certain kinds of constraints. For example, we initially have screened for anomalies without regard to ramp rate constraints and manually verified that units were ramp rate constrained. This process served to test the accuracy of the indicator for ramp constrained units.

8 This condition will hold unless a load balance constraint or ancillary service requirement is not satisfied, in which case all resources could be at their upper or lower limits. We understand that at present, prices should be set in these situations by the highest cost accepted offer. Units with binding energy limits are on the margin if

this location should be equal to the energy offer price of the marginal generator or the bid of the marginal price capped load at its dispatch level plus any relevant opportunity costs.[9] Similarly, in an interval with n binding transmission constraints, at least n+1 generators or price capped load bids should be on the margin.[10] The marginal generators and load bids are identified in the pricing dispatch in which some resources have schedules fixed in the physical dispatch and are not eligible to set price and in which fixed block units such as gas and combustion turbines are treated as if they are dispatchable at any level between zero and their upper operating limit. If the number of marginal units is less than the number of binding transmission constraints plus one, this is reported by the price validation tool for further review.

Because of the complexity and number of the potential tradeoffs between energy and ancillary service schedules and intertemporal tradeoffs, units that are marginal due to these trade-offs are not directly identified by the price validation software. Trade-offs between energy and ancillary service schedules are validated by verifying that the energy dispatch and ancillary service schedules correctly reflect the opportunity cost of energy ancillary service tradeoffs.

3. Evaluate Energy Schedules for Consistency with LMP Energy Prices

This test reviews the energy schedules for generating units and loads submitting price capped bids in the day-ahead market, real-time pre-dispatch and real-time dispatch and verifies that they are consistent with the LMP prices at those locations. For the purpose of

---

the sum of their offer price and the shadow price of the energy constraint is equal to the LMP price at their location. We have applied tests to verify that the daily schedules of energy limited resources do not exceed their maximum energy limit nor fall below the minimum energy limit and to verify that the constraint shadow price is correctly reflected in prices when the energy limits are binding.

9 In the IFM and real-time pre-dispatch, energy prices can reflect the opportunity cost of providing energy instead of ancillary services and can also reflect the opportunity cost of using energy limited resources with binding daily energy limits. Energy prices can also be set by intertemporal opportunity costs on ramp-constrained units in IFM, real-time pre-dispatch and the real-time dispatch. For example, in one RTPD case, prices were high in one 5-minute interval then low in the following four intervals. One unit is ramp-constrained down throughout the first three low-priced intervals, so if it is dispatched up an additional megawatt in the high priced interval, its output will be 1 megawatt higher in each of the next three intervals in which it would be operating uneconomically. The resource's opportunity cost in the high priced interval is therefore equal to its losses in the following three periods in which it operates uneconomically and this opportunity cost sets the incremental cost of meeting load at some locations during the high-priced interval.

10 This condition should hold unless a transmission constraint cannot be solved in which case all resources could be at their upper or lower limits on one or both "sides" of the constraint. Constraint shadow prices would then be determined by constraint violation costs and energy and ancillary services prices could be determined in part by these constraint violation costs. We have applied tests to verify that constraint shadow prices are set to the appropriate value when a transmission constraint cannot be solved. Units with binding energy limits are on the margin if the sum of their offer price and the shadow price of the energy constraint is equal to the LMP price at their location.

this test, generating units and price capped load bids are divided into five categories based on their schedules: (1) generating units or load bids that were on the margin;[11] (2) generating units that were scheduled to operate at their upper limit (given their ancillary service schedules) or were ramp constrained up; (3) price capped load bids that were dispatched down to their lower limits; (4) generating units that were scheduled to operate at their lower limit (given their ancillary service schedules) or were ramp constrained down; and (5) price capped load bids that were dispatched to their upper limit.

The next step is to verify that: (1) the LMP at the location of each marginal resource or price capped load bid was equal to the resource or load's energy bid at its dispatch point (the same as test 2);[12] (2) the LMP at the location of each resource scheduled at its upper limit or that was ramp-constrained up was greater than or equal to the unit's energy bid at its dispatch point; (3) the LMP at the location of each price capped load scheduled to its lower limit was greater than or equal to the load's bid at its dispatch point; (4) the LMP at the location of each unit scheduled at its lower limit or that was ramp-constrained down was less than or equal to the unit's energy bid at its dispatch point; and (5) the LMP at the location of each priced capped load scheduled at its upper limit was less than or equal to the load's bid at its dispatch point.

4. Replicate Ancillary Service Prices from Shadow Prices

This test verifies that ancillary prices could be replicated from appropriately calculated shadow prices. The application of this test evolved in the course of testing with changes in the number and nesting structure of ancillary service regions and with changes in the way ancillary service shadow prices are reported.

5. Validate Ancillary Service Prices Based on Marginal Suppliers

For the day-ahead market and real-time pre-dispatch, this test identifies the marginal ancillary service suppliers whose offers determined the market-clearing prices for spinning reserves, 10-minute reserves, regulation up and regulation down, in each ancillary service region in each hour, and validates the cascaded market clearing prices for each ancillary service. The offer of the marginal ancillary service supplier may have an energy market opportunity cost that sets the ancillary service price.

---

[11] This could include generation or load self-schedules that were curtailed at the penalty value in the pricing pass of the IFM or RTUC software.

[12] In the case of energy limited units with binding energy constraints, we verify that the sum of their offer price and the shadow price of the energy constraint is equal to the LMP price at their location.

6. Examine Ancillary Service Schedules for Consistency with Prices

For the day-ahead market and real-time pre-dispatch, this test verifies that the ancillary service schedules for ancillary service suppliers were consistent with the market clearing prices for energy, spinning reserves and 10-minute reserves, regulation up and regulation down in each region and in each hour or pre-dispatch period. This evaluation includes verification that the sum of the ancillary service offer price and energy market opportunity cost is less than or equal to the price of ancillary services for all resources scheduled to provide each ancillary service.

7. Examine Unit Commitment for Consistency with Prices

The day-ahead market unit commitment results are validated by applying two tests to generating unit schedules. First, we determine whether there were units that were not committed that could have profitably operated at the day-ahead LMPs, taking into account the resource's minimum load cost and start-up costs over the commitment period. Second, units whose schedules would require uplift payments because their day-ahead energy and ancillary service revenue is less than their as bid costs are identified. These tests were also applied, but less comprehensively, to the real-time unit commitment decisions by RTUC.

Because of the non-convexities inherent in unit commitment decisions, not all anomalies identified by these tests indicate software flaws or limitations. Our review focused on identifying material anomalies or patterns involving a large number of smaller anomalies within a single test case. We did not attempt to definitively resolve individual small discrepancies within a given test case that would be consistent with small changes in loss factors or prices associated with changes in the unit commitment.

B. Scope of Testing

This report covers the testing of certain elements of the CAISO market software that we have carried out. It does not cover all testing that has been carried out by the CAISO. In particular, this report does not cover the performance of various associated systems. For the purpose of the testing covered in this report:

- The IFM Cases analyzed have not been run in conjunction with RUC or market power mitigation.[13]
- Some of the real-time cases were taken from the market simulation process and the input data went through SIBR and the market power mitigation steps.

---

[13] The CAISO is independently testing RUC and the market power mitigation process.

Other real-time cases were run in a test environment without market power mitigation, without going through SIBR, and without real-time events such as operator actions and generator and transmission outages.

- The RTUC and RTD cases have been tested independently of the IFM cases so we have not tested issues relating to the relationship between IFM, RTUC and RTD schedules.

- We have not tested associated processes such as the load forecaster.

The results of testing these other elements of the CAISO market software will be covered elsewhere.

C. Test Cases Analyzed

In the third stage of analysis track testing, LECG has fully reviewed 11 IFM cases:

(1) base case .5% and .01% MIP gap

(2) high load base case fc_ifm_patch190_case1c_highdemand_S617263182

(3) fc_ifm_ a2 test7a_e032408_ s917271260 (referred to as case 7a below)[14]

(4) fc_ifm_t3_test108_ e032208_s617263382 (referred to as case 108 below)[15]

(5) fc_ifm_t3_case104_p197_ e032808_S617263462 (referred to case 104 below)

(6) fc_ifm_t3_test102_e032808_ S617263443 (referred to as case 102 below)

(7) case fc_ifm_a2_case106_ e032408_s917271280 (referred to as case 106)

(8) case fc_ifm_a2_case107_ e032408_s917271240 (referred to as case 107)

(9) fc_ifm_t3_test109_patch231_ e050608_s617264882 (referred to as case 109)

(10) fc_ifm_t3_test110_patch231_ e042308_s61726484 (referred to as case 110) and

(11) fc_ifm_t3_test111_patch228_ e042308_s617264602 (referred to as case 111).

Three of these cases were rerun on patch 284 in June to verify that previously identified problems had been corrected. These were case 7a, case 110 and the base

---

[14] This is also sometimes described as case 105 in recent CAISO case documentation. We have continued to refer to it as case 7a, as in our Preliminary Report.

[15] This case was initially run with a .5% MIP gap and was subsequently rerun with a .01% MIP gap.

case.[16] Cases 7a and 110 were completely reanalyzed, the base case was only reviewed to confirm that the changes in the definition of the objective function caused the solution to converge to the expected unit commitment solution within the intended MIP gap specification.[17]

Since the Interim report dated July 1, 2008, LECG has received and fully analyzed an additional ten IFM cases:

(1) ifm_test_p324_e071608_mktsimanalysis_517269024 (referred to as the marketsim case below)

(2) fc_ifm_t3_wheeling_p324_e072408_617266263 (referred to as the wheeling1 case below)

(3) fc_ifm_t3_wheeling_p324_e072408_617266323 (referred to as the wheeling2 case below)

(4) fc_ifm_t3_test110_patch321_e071308_s617266082 (referred to as case 110-3 below)

(5) fc_ifm_t3_highdemand_fz_p323_e071608_s617266163 (referred to as the forbidden zone case below)

(6) fc_ifm_t3_test110_patch360_e081508_s617266462 (referred to as case 110-4 below)

(7) fcreg_ifm_t3_case105_var5990_validation_p363_e082108_617266763 (referred to as case 7a-3 below)

(8) fc_ifm_test_block_bid_p20031_e091108_s517280961 (referred to as block2 case below). This case is a rerun of the original block transaction case (C in the section below on cases that were not fully analyzed).

(9) fc_ifm_staging_fbzcase_patch20040_e091408_s417209455 (referred to below as forbidden2 case.)

(10) fc_ifm_staging_fbzcase_patch20050_e092308_s417209936 (referred to below as forbidden3 case.)

---

16 The rerun of case 7a is called case 105 in CAISO documentation. The rerun is not exactly the same as the original case 7a, which could not be reloaded because of the number of intervening changes. The rerun 7a uses a more recent database but tests the same issues as the original case 7a.

17 Cases will not necessarily solve to the specified MIP gap, as they may time-out before achieving that value.

In addition, we have received three other cases for which a full analysis has not been carried out:

A) ifm_e072408_217266622_mktsim_contingencies – This case was run so that we could test the recalculation of prices in a case where contingency constraints were binding.

B) fc_ifm_test_nomo_p315_e070908_517268664 (referred to as the nomogram case below) – This case was provided to test prices in a case where nomogram constraints were binding. We only looked at the price recalculation in this case.

C) fc_ifm_test_block_bid_e062908_s517268343 – This case was generated to provide an example of a block transaction. We did not do a complete evaluation of this case.

D) fc_ifm_test_block_bid_p20031_e091108_s517280961 – This case is a re-run of the block2 case that was provided to determine if a software patch fixed a dispatch error at a pump storage unit in the block2 case.

We also reviewed portions of four additional cases:

A) fc_ifm_a2_case2001_patch208_e917271501_rerun and fc_ifm_t3_patch231_ e050608_lap_s617265042. These cases were only reviewed to validate the prices calculated for aggregate nodes.

B) fc_ifm_t3_case4001_p228_e042308_s617264003 and fc_ifm_t3_case4001_p228_ e042308_s617264804. These cases were used for verifying the "Distributed Generation Slack" functionality vs. "Distributed Load Slack." We were asked to verify that we could replicate the prices reported for these cases.

We have to date reviewed 17 RTUC cases:

(1) fc_rtpd_t3_patch208_ e302808_ mipgap.01% and mipgap .5% (referred to as the base case)[18]

(2) fc_rtpd_t3_test201_patch208_e032808_s617282281 (referred to as case 201)

(3) case fc_rtpd_t3_case203_patch208_e_032808_s617282321 (referred to as case 203)

(4) case fc_rtpd_t3_case204_p208_e032808_s617282324 (referred to as case 204)

[18] This case is also described as case 202 in some CAISO documentation.

(5) case fc_rtpd_t3_case205_patch208_e_032808_s617282323 (referred to as case 205)

(6) case fc_rtpd_t3_case206_patch208_e_032808_s617282341 (referred to as Case 206)

(7) case fc_rtpd_t3_test207_patch208_e032808_s617282322 (referred to as Case 207)

(8) case fc_rtpd_t3_test208(1)_patch228_e042308_s617282621 (referred to as Case 208(1))

(9) fc_rtpd_t3_test208(2)_patch228_e042308_s617282622 (referred to as Case 208(2))

(10) fc_rtpd_t3_test209_patch231_e050608_s617282801 (referred to as Case 209)

(11) fc_rtpd_case210_patch262_e052308_s617282981 (referred to as Case 210) and

(12) fc_rtpd_test_case206_p353_e080108_s519043870 (referred to as Case 206b).

(13) fc_rtpd_staging_export_binding_patch20058_e100208_s427291852 (referred to as export case below)

(14) fc_rtpd_u2_export_binding_patch20058_e100208_s217487268 (referred to as export2 case below)

(15) fc_rtpd_staging_lap_patch20058_e100208_s427295501 (referred to as lap3 case below)

(16) fc-rtpd-u2-stuc-patch20065-e100208-s217493325-rerun427291852 (referred to as stuc case below)

(17) fc-rtpd-u2-lap-patch20065-e100208-s217493326-rerun429295501 referred to as u2_lap case below)

Three of these cases were rerun, two on patch 283 and one on patch 284, in June to verify that previously identified problems had been corrected.[19] These were cases 202, 203 and 204.

We also reviewed one additional case, fc_rtpd_t3_lap_p283e061508 s617283049. The purpose of this case was to provide a complete set of data for testing the recalculation of Apnode pricing for RTUC. Finally, we reviewed fc_rtpd_test_hourly_ dispatch_e060908_s519014598 (referred to as hourly dispatch case) to evaluate whether a multi-hour block transaction was scheduled correctly.

---

[19] Patch 283 includes most changes but did not include the reformulated objective function.

During May and June 2008 we reviewed four RTD cases:

(1) fc_rtd_t3_basecase_patch269_e060108_s617283024 (referred to as the Base Case)

(2) fc_rtd_t3_test301_patch269_e060108_s617283028 (referred to as Case 301)

(3) fc_rtd_t3_test303_patch273_e060108_s617283035 (referred to as Case 303) and

(4) fc_rtd_t3_test304_patch273_e060108_s617283033 (referred to as Case 304).

We reviewed case fc_rtd_t3_test301_patch284_e061508_ s617283048, which provided a complete set of data for testing the recalculation of Apnode pricing for RTD.

We also reviewed three additional cases in order to validate LAP prices: fc_rtd_t3_test301_patch297_e062408_s617283068 (referred to as case 301 LAP), fc_rtd_t3_lap_patch297_e062408_s617283074 (referred to as patch 297 LAP), and fc_rtd_t3_lap_patch324_e072208_617283127 (referred to as patch 324 LAP).

The test cases we have reviewed to date have a number of features that have helped verify that particular features of the CAISO software are working as intended. The conditions we have observed in testing, and in which the pricing software operated correctly, include:

- Unsolved internal transmission constraints with shadow prices equal to the constraint violation penalty.
- Binding external tie-line scheduling constraints with shadow prices as high as $2,000 in RTUC cases during the scheduling and price setting hour.[20]
- Regional ancillary service prices reflecting minimum ancillary service requirements.
- Price-capped load bids[21] setting prices in the day-ahead market.

[20] We observed tie-line constraint shadow prices as high as $30,000 in a number of RTUC cases during the first few intervals of the case in which tie-line schedules are fixed at the level in the prior hour. These high shadow prices arise because the tie limits are lower than those enforced in setting schedules for those intervals and the intertie constraints were violated by the fixed schedules. Such $30,000 constraint shadow prices were never observed in the second hour of the RTUC case in which the tie-line schedules were determined, precisely because tie-line schedules could be adjusted to avoid violating the constraints.

[21] We use the term "price capped load bids" to refer to bids to buy power in the day-ahead market that are conditioned on the price at the specified location being less than or equal to the price specified by the buyer (i.e., less than or equal to the buyer's bid price).

- Binding intertemporal limits on energy limited resources in the day-ahead market.
- Reserve shortages within particular ancillary service subregions in the day-ahead market, real-time pre-dispatch or real-time dispatch.
- CAISO-wide reserve shortages in the day-ahead market, real-time pre-dispatch or real-time dispatch.
- Changes in transmission limits over the analysis period.
- Exports curtailed at the price cap in the day-ahead market.
- Self-scheduled load that could not be met.
- Forbidden region constraints enforced.
- Wheeling transactions present.
- Multi-hour block transactions offered and scheduled.
- Binding Nomogram and Contingency Constraints

Not all elements of the market software have been tested by the conditions included in the test cases reviewed for this report. Some of the conditions not verified in this report include:

- No RTUC test cases have included uneconomic COG units running due to minimum run time constraints.

D. Test Results

1. Replicate LMP Energy Prices

IFM

<u>Pnode LMP Replication</u>

We are able to replicate the congestion component and overall LMP price for all prices reported by the IFM software in the cases tested. The issues that were identified in the preliminary and interim reports relating to the prices reported for disconnected units, the prices in some cases with very high constraint shadow prices, or on export constrained interties have been addressed and were not observed in cases run or rerun on the latest software patches. The only remaining anomalies appear to arise from master file network model issues. In addition, we were able to replicate prices in the 4001 cases used for testing the distributed generation and load slack buses. Until recent changes, our replication of LMP prices was based on shift factors rounded to six decimal places and loss factors rounded to four decimal places to be consistent with the degree of precision in the CAISO market software. Recent changes to the software have resulted in shift factors being rounded to five decimal places. Loss factors are still rounded to four decimal places. The resolution of price calculation issues is reviewed below.

We were initially unable to replicate the congestion component of prices reported by the IFM software for seven disconnected Pnode IDs in IFM cases 7a, 106 and 107 using shift factors, offsets, and transmission constraint shadow prices reported by the IFM software. The prices we calculated differed from the prices reported by the IFM software by as little as a few cents or as much as $48/MWh in case 7a, as much as $95/MWh in case 106, and as much as $523/MWh in case 107. We observed price calculation errors on the same seven Pnodes in the base test case run with patch 157. The message files for these cases indicated that these seven Pnodes were electrically disconnected from the grid in these cases, while connected in the remaining cases in which the Pnode prices for these locations could be replicated. It appears that we were unable to replicate the prices from the shift factors and offsets because the shift factors exported from the IFM for these seven locations were non-zero, while the prices reported by the IFM software for these locations were calculated with a zero shift factor (presumably reflecting the disconnection from the grid) and a non-zero offset.[22] Case 7a was rerun on the latest patch in June, none of these units were reported to be disconnected from the grid, and there were no price calculation errors.

We were also initially unable to replicate the congestion component of prices reported by the IFM software for IFM case 110. This inability appears to be a result of inconsistent rounding of shift factors in the calculation of prices. Siemens indicated that resources that received non-zero schedules in the IFM have their shift factors rounded to six decimal places before calculating prices, while resources that did not receive an energy schedule did not have their shift factors rounded. As a result, we were able to replicate prices for all resources using either the rounded or unrounded shift factors, but not using any single rule. Changes were made to address these inconsistencies and when we evaluated a rerun of IFM case 110, almost all of these discrepancies were eliminated, the number of discrepancies falling from 20,710 to 5. The five remaining anomalies appear to be due to inconsistent shift factor rounding for shift factor values between .00001 and .000001. The shift factor dropping issue for shift factors between .00001 and

---

[22] The penalty factor for these locations is non-zero and we are able to replicate the loss component of these prices using that penalty factor, so these locations do not appear to be disconnected from the grid in determining the loss factor.

.000001 was also seen in the wheeling1 case, wheeling2 case, forbidden region case, and case 100-3. Siemens began rounding to the precision of the SF_THRESHOLD variable, currently set to .00001 in patch 360. Since this change was implemented, we have been able to replicate all congestion components.

We initially observed a problem with the reporting of LMP prices for Pnodes on export constrained interties in the marketsim case and the nomogram case. The shadow price on binding interties was being added with the wrong sign to the prices reported for these locations. This was a price calculation error; the dispatch was correct and consistent with the correctly calculated price. The software issue was corrected and CAISO re-ran type B case 2006 which had an export constraint binding. We have confirmed that the export constrained intertie prices are calculated correctly with the correct sign.

We initially observed an issue in the forbidden2 case in which one Pnode was reported as having congestion in an hour during which there were no binding nomogram or flowgate constraints. This Pnode was located on an inter-tie, however, that inter-tie was not binding and no other inter-tie was binding with a shadow price equal to this Pnode's congestion component. This problem was not present in the forbidden3 case run on a later patch.

We initially observed an apparent data issue in the rerun case 7a in which one Pnode was linked to two interties. CAISO has verified that this issue has been fixed in the master file, but cases 7 and 7a-3 are re-runs of an old case using old master file data. In this case some resources on the Pnode were linked to one intertie while others were linked to another intertie. This resulted in inconsistent prices for resources at the same Pnode when one intertie was binding and the other was not. A similar mapping issue was seen in the marketsim case where two resources on the same Pnode were not mapped to the same intertie ID. CAISO believes this issue is arising from a data mapping problem that will be resolved in a master file update. A similar instance was seen in case 7a-3.

<u>Apnode LMP Replication</u>

We were initially unable to recalculate a variety of Apnode prices in IFM test cases. This problem was eventually traced to incorrect nodal weighting data being exported from the IFM program. Case 105[23] was rerun with changes to ensure that the correct nodal load weights were saved and exported and we were able to exactly replicate all Apnode prices, including the default LAP prices. This process is reviewed below.

[23] fcreg_ifm_t3_case105_p297_e062408_617265884

We were initially unable to recalculate certain Apnode LMPs in the base cases, cases 7a, 102, 104, 106, 107,108, 109, and 111 due to missing data for the relevant load weights. The affected Apnode LMPs were those for LAPs and LAFs.[24] The rerun IFM case 110 was also missing data for the relevant load weights. The case 110-3, wheeling1 case, wheeling2 case, and forbidden region case were also missing data for the relevant load weights. This is a master file data issue not related to the pricing software.

We were initially unable to calculate the Apnode price with static participation factors at one Apnode in the marketsim case, wheeling1 case, and wheeling2 case because they had component Pnodes that were affected by the price recalculation errors described in the Pnode LMP Replication section above relating to the dropping of small shift factors or the incorrect reporting of LMPs on Pnodes located on export constrained interties. These issues have been resolved with the changes to rounding precision discussed above.

The CAISO provided us with data from IFM case 2001 which included the load weights required to replicate Apnode prices. We recalculated the SCE, PG&E and SDG&E LAP prices from the component Pnode prices and were able to replicate the SCE and SDG&E LAP prices. We were not able to replicate the PG&E LAP price in any hour. The differences ranged from a penny up to about $.75/MWh. We observed the same inability to replicate PG&E prices in case fc_ifm_t3_patch231_ e050608_lap_s617265042. We were not able to carry out replication of all of the aggregated generator nodes for case 2001 because the relevant Apnode prices were not correctly exported, resulting in the export of zero prices for Apnodes that clearly had non-zero prices. We were able to replicate the trading hub prices for case 2001. We also observed Apnode price recalculation errors in case 110 for aggregate generators, loads, and LAPs. The recalculated prices differed from the prices reported in the software by several cents. All off these issues were traced to the export of incorrect load weight data as noted above. This was corrected and we have been able to correctly calculate these Apnode prices in all subsequent cases.

LECG is unable to replicate certain Apnodes prices with types Aggregate System Resource and Aggregate Generator. The prices for these Apnodes are reported as blanks in pricing data files. CAISO has determined that this is due to a master file issue not related to the pricing software in which some Apnodes have multiple IDs in the master file and only one is being populated with prices.

As noted above, all of the apparent price calculation errors for aggregate nodes appear to have simply reflected the export of incorrect nodal load weights and we have

---

[24] The emm_scuc_imm_laf.csv and emm_scuc_imm_lap.csv files for the IFM test cases other than case 2001 are blank.

been able to recalculate every Apnode price in the rerun of case 105 in which the correct weights were exported.

RTUC

Pnode LMP Replication

We are able to replicate the reported Pnode LMP prices at almost every location in every rerun test case. The few remaining inconsistencies appear to arise from master file network model issues.

When we initially carried out the replication of Pnode prices for a number of RTUC cases, we were able to replicate the prices at some Pnodes using unrounded shift factors and to replicate Pnode prices at other Pnodes using shift factors rounded to 6 decimal places. The issue was caused by the inconsistent dropping of shift factors between .00001 and .000001. As discussed in the IFM section above, Siemens changed their rounding of shift factors from six decimal places to five decimal places and we were able to verify all congestion components in the IFM cases. The export, export2, and lap3 cases that were processed after this change was added to the code confirmed that this issue has also been fixed in RTUC.

We initially observed numerous instances in which we were unable to recalculate the loss component of the LMP price in cases 208(1) and 208(2). After the CAISO determined that loss penalty factors had been rounded in the calculation of the LMP, we revised our tests to account for the specified rounding and are now able to recalculate the loss components as reported by the RTUC software. We did not observe any loss component recalculation issues in the rerun of the RTUC basecase or in cases 203 or 204.

We observed two instances in the rerun of the RTUC basecase and one instance in the rerun of RTUC case 204 in which we were unable to recalculate the congestion component of the LMP at a Pnode location located on an intertie. The reported price differs from our calculated price by over $200 in some cases. This Pnode is only linked to one intertie and there is no binding intertie with a shadow price equal to the price difference between calculated and reported, therefore this does not appear to be related to an intertie mapping issue. The CAISO was going to write a variance on this issue to correct a master file network model issue.

We initially observed the same problem with the reporting of LMP prices for Pnodes on export constrained interties in the hourly dispatch case that was identified in the IFM marketsim and nomogram cases. The shadow price on the binding interties is of the wrong sign in the prices reported for these locations. As in IFM, this was a price calculation error; the dispatch was correct and consistent with the correctly calculated price. As discussed above, we have rerun IFM cases to verify that export prices are correctly calculated in IFM. We have also tested this in the RTUC export and export2 cases, and verified that the export shadow price was correctly applied in calculating the intertie prices.

Apnode LMP Replication

We were initially unable to recalculate a variety of Apnode prices in RTUC test cases. As in IFM, this problem was eventually traced to incorrect nodal weighting data being exported from the RTUC program. The RTUC lap case[25] was rerun with changes to verify that the correct nodal load weights were saved and exported and we were able to exactly replicate all Apnode prices.

We were initially unable to recalculate Apnode LMPs for a number of Apnodes in the RTUC base case and case 206b that were composed of Pnodes whose prices cannot be recalculated as noted above. In addition, we were unable to recalculate the LAP prices for the RTUC basecase due to a lack of mapping data.[26] In the RTUC Case 210, mapping data was provided, but certain Apnode prices could not be recalculated from the underlying Pnode prices, including the PG&E default LAP price.[27] These discrepancies were due the export of incorrect nodal load weights. With the correct nodal load weights exported, we were able to recalculate all Apnode prices in the export2 case and the rerun lap3 case.

CAISO also provided case fc_rtpd_t3_lap_p283_e061508_s617283049 for the purpose of validating Apnode prices. We identified 28 instances in which we could not replicate the prices at aggregate generator or LAP locations. These price recalculation errors ranged from $0.01 to $19.43/MWh. In addition, we calculated a non-zero price at two Apnodes in this case for which the CAISO software reported a $0 price. This issue was resolved when the correct nodal load weights were exported, as noted above.

We are unable to replicate certain Apnodes prices with types Aggregate System Resource and Aggregate Generator. The prices for these Apnodes are reported as blanks in pricing data files. CAISO has determined that this is due to a master file issue in which some Apnodes have multiple IDs in the master file and only one is being populated with prices.

---

25 fc_rtpd_t3_lap_closeuconn_p297_e062408_s617283070

26 We were unable to recalculate the Apnode prices for numerous aggregate load and LAP locations in rerun RTUC cases 202, 203, and 204 because of missing data for the relevant load weights.

27 The discrepancies ranged from a few cents up to $1.66 per MWh.

We were unable to recalculate some Aggregate Generator prices in the export case. This is related to the incomplete export of dynamic participation factors and does not reflect an incorrect calculation of prices.

As noted above, all of these apparent price calculation errors for aggregate nodes appear to have simply reflected the export of incorrect nodal load weights and we have been able to recalculate every Apnode price in the rerun of the LAP test case in which the correct weights were exported.

RTD

Pnode LMP Replication

All Pnode LMP prices could be recalculated except for two Pnode LMP prices in the RTD base case. These prices could not be recalculated because the Pnodes are mapped to two different interties having different congestion components. LECG is able to validate the reported LMP price using one of the constraint shadow prices but not the other. The CAISO is aware of the apparent mapping problem and is working to correct it. All other Pnode LMP prices in the other RTD cases could be recalculated.

Apnode LMP Replication

We attempted to recalculate Apnode LMP prices in cases 301, 303 and 304. In each case, there were a handful of Apnodes for which we initially could not validate the LMP as a result of nodal load weight data issues.

CAISO also provided case fc_rtd_t3_test301_patch284_e061508_s617283048 for the purpose of validating Apnode prices. LECG observed 28 instances where we could not replicate the prices at aggregate generator or LAP locations. These price recalculation inconsistencies ranged from $0.04 to $474.00. After the RTD software was modified to export the correct load weights, wee examined RTID case patch 324 LAP in order to validate LAP prices. We were able to correctly recalculate the all Apnode prices, including the LAP prices in this case.

2. Validate LMP Energy Prices based on Marginal Offers/Constraint Violation Costs

Analysis of marginal units determining LMP prices was not carried out for the November cases, but was included in the evaluation of the April and June 2008 analysis track cases.

IFM

The number of marginal units appropriately exceeded the number of binding transmission constraints in every hour of the base cases, and analysis track cases 7a (rerun), 102 104, 106, 107, and 111, and in all hours of case 108 except hours 22, 23, 0 and 3 (GMT), and all but one hour of cases 109 and the original case 7a. The appropriate number of marginal units also appears likely to exist in the remaining three hours of case 108 and one hour of case 109; however, because of the shift factor rounding issues discussed in Subsection 3 below, there were a number of units which should be marginal whose offer prices appear to differ by a penny or so from the LMP price at their location. There is one hour of the original case 7a in which curtailed generator self-schedules set prices in the scheduling pass, and we were not able to identify one of the tradeoffs determining prices in the pricing pass. The number of marginal resources appropriately exceeds the number of binding transmission constraints in every hour of case 110-3, 7a-3, the block2 case, forbidden2 case, and forbidden3 case.

In case 108 the 30970_MIDWAY_230_30973_SUNST _230_BR_1 _1 flowgate is violated and appears in the emm_scuc_output_flowgate_v.csv file with a shadow price equal to $3,000 per megawatt (either $1,000 or $3,000 in cases 109 and 110). This constraint violation cost correctly sets prices in the relevant hours.

There are also some flowgates that are overloaded in the pricing pass, and have shadow prices in excess of pricing pass constraint violation penalty ($3,000) in cases 109 and 110. This outcome is a result of the way constraints are relaxed in the pricing pass. Transmission constraints that are violated in the scheduling pass are relaxed in the pricing pass at the specified constraint violation cost for a small range beyond their scheduling pass value. If the violated constraint were the only such constraint in the scheduling pass, one would expect that the constraint would be violated to the same or lesser extent in the pricing pass and the constraint violation penalty would set prices.[28] If there is more than one constraint violated in the scheduling pass, this intuition may not hold in the pricing pass solution, and constraints can bind in the pricing run at values well in excess of the pricing run constraint violation cost. This was not observed in rerun IFM cases 7a and 110 because pricing run constraint violation costs were set at the same level as in the scheduling run.

The observed outcomes with some transmission constraint shadow prices above the pricing pass constraint relaxation penalty price are consistent with the intended software implementation of transmission constraint relaxation in the scheduling and pricing passes. Because constraint shadow prices in the pricing run can exceed the pricing run constraint violation cost, it is possible for prices to be set by constraint

---

[28] If the incremental cost of the resources dispatched to solve the constraint in the scheduling pass was greater than the constraint violation cost, one should expect the constraint to be binding in the pricing pass at a price in excess of the pricing pass constraint violation cost. If the incremental cost of the resources dispatched in the scheduling pass was less than the pricing pass constraint violation cost, then one would expect the price in the pricing pass to be set by the pricing pass constraint violation cost.

shadow prices that are in excess of the pricing run constraint violation cost, up to the value of the scheduling run constraint violation cost. In these cases, prices are not set by the constraint violation costs but by the bids and offers of marginal buyers or sellers and assure that the LMP prices recover the actual cost of meeting load or accommodating transmission schedules. In wheeling1 case there was one instance in which the shadow cost on a violated constraint was greater than the high pricing run penalty and was also higher than the scheduling run penalty. This outcome was likely related to the program timing out before an optimal solution was found. This did not occur in the wheeling2 case.

In cases 108, 109 and 110 some flowgates are not violated, but are at their limit, and also have shadow prices in excess of the pricing run constraint violation cost. For example in case 108, the METCALF_MORGANHL_ BG flowgate in HB 15 PST is binding with a shadow price of $12,993.90 in the pricing pass.[29] As explained above, in the current design constraints that are not violated in the scheduling pass are not eligible to be relaxed at the constraint violation cost in the pricing pass (i.e., they are hard constraints in the pricing pass) and this accounts for the observed outcome. This was also observed in rerun IFM case 110.

RTUC

The number of marginal units plus violated transmission constraints appropriately exceeded the number of binding transmission constraints in all of the RTUC cases.[30] As in IFM case 108, there were transmission constraints that could not be solved and were violated in the RTUC base case. The same kind of outcomes that were identified in IFM case 108 were present in this base case, in particular, there were a number of constraints that were violated or binding in the pricing pass with shadow prices in excess of $1,000 (the constraint violation cost for constraints relaxed in the pricing pass). We believe these outcomes are the expected result of the way constraints are currently relaxed in the pricing pass. These outcomes for violated constraints were also present in the rerun of the RTUC basecase and in cases 203 and 204.

---

[29] The analysis track cases included cases with varying constraint violation penalties, including some with higher penalties than currently envisioned for MRTU implementation. Later analysis track cases had penalties set so that constraint shadow prices did not exceed $5,000.

[30] Marginal units were identified for the second hour of the RTUC case which sets interchange and ancillary service prices. The dispatch consistency test, discussed in Section 3, was applied to all intervals of the RTUC cases.

RTD

The correct number of marginal units has been identified in the RTD basecase and cases 301, 303 and 304.

3. Examine Energy Prices for Consistency with Energy Dispatch

IFM

We have verified that in the cases we reviewed that resources are being dispatched consistent with their bids, the LMP prices at their location, and the CAISO market design in almost all instances, and none of the unresolved inconsistencies indicates the existence of a flaw in the calculation of prices. The process of resolving issues is reviewed below.

In case 108, 110, and 111 we initially observed instances in which generating units are dispatched to a point on their bid curve that is uneconomic by more than \$0.01, but less than \$0.02.[31] The threshold that LECG, the CAISO and Siemens have agreed should be applied to distinguish dispatch and pricing issues from the effect of rounding conventions is \$.01, so these errors, while small, are outside the bounds we have used to identify software issues for review. The very high transmission constraint shadow prices in this case raised the possibility that this discrepancy was a result of a difference in shift factor rounding within the IFM software that only becomes apparent with such high shadow prices. As explained above, the shift factors used for the optimization are being rounded to six decimal places while our replication of prices was originally based on unrounded shift factors. This rounding does not account for the dispatch inconsistencies in these cases, however, as the prices we use for the comparison are based on the rounded shift factors used in the optimization and the rounding in these cases only changes prices by a fraction of a cent. These issues were addressed by Siemens and were not present in the rerun IFM cases 7a or 110.

There were no load schedule anomalies in the base cases or cases 102, 104, 106, 107, or 111. Price capped load was also correctly scheduled in both the scheduling pass and the pricing pass of the rerun of case 7a. Price capped load was also correctly scheduled in both the scheduling pass and the pricing pass of case 110-3, wheeling2 case, and the forbidden region case.

We initially observed a price-capped load bid in case 7a that was scheduled to a level that was inconsistent with its bid and the price at its location. This load had a bid price of \$5/MWh and a LMP of \$0.00 at its location, but the schedule in the pricing pass

[31] These price and dispatch inconsistencies were present in both the .5 and .01% MIP gap runs for case 108.

was 0.033 MWs below the unit's bid MW value, which should have been fully dispatched. The dispatch inconsistency was very small, but it should not exist and it did not appear to be related to any of the other issues that had been identified. The load bid was scheduled to the appropriate level in the scheduling pass. The cause of this anomaly was never identified and it is not clear which software change eliminated the anomaly, but as noted above it was not present in the rerun of case 7a.

We identified one instance in the initial run of case 110 of a unit turning on and ramping up to its lower limit but no further during the hour, despite having an incremental cost that was lower than the price at its location and having sufficient ramp capability to reach a higher output.[32] This unit was ramped correctly in the rerun of case 110.

We also initially observed instances in cases 108, 109, and 110 in which load bids were dispatched to a place on their bid curve that was uneconomic by more than $0.01, but less than $0.02. This pattern was likely related to the observation regarding generation units and shift factor rounding noted above and these inconsistencies were not present in the rerun of cases 7a or 110.

There were no wheeling transactions in the base cases, or in analysis track cases 7a, 102, 104, 106, 107, 108, 109, 110, 111, or the rerun of cases 7a or 110, so no wheeling transaction anomalies were identified. There were also no wheeling transactions in case 110-3 or the forbidden region case. There were wheeling transactions in the wheeling1 case, wheeling2 case, and marketsim case. In all instances, valid wheeling transactions were scheduled correctly. These cases contained some invalid wheeling transactions, unbalanced wheels, because the data used to create the cases was not entered through SIBR.

We identified a data issue in the base cases and cases 7a, 102, 104, 106, 107,108, 109, 110, and 111 in which units are committed for regulation and are ramp constrained using their regulation ramp rate, but are not flagged in the data as ramp constrained. Since the original report, it has been determined that the upper and lower limits in another table (emm_scuc_output_bid) can also be used to identify units that are not dispatched marginally because of limits, but use of this table does not resolve all of the omissions, which were also seen in rerun IFM cases 7a and 110. This was also seen in the wheeling1 case, wheeling2 case, marketsim case, forbidden region case, case 110-3, forbidden2 case, and forbidden3 case. This issue does not reflect any price calculation

---

[32] Our understanding of the prescribed start-up ramp rule is that a unit can ramp during the hour in which it comes on-line to its lower limit plus one-half of its ramp rate times, either 20 minutes or 60 minutes, depending on whether it's a fast or slow ramping unit.

error or dispatch issue but should be corrected prior to market implementation to speed price verification.

In rerun IFM case 7a, we initially observed resources on interties that were dispatched as though they were marginal within their self-scheduled region although the LMP at their location was between $0.10 and $0.30 greater than the self-schedule penalty price. We also observed two resources on interties that had energy schedules that were inconsistent with their bids and prices. These units were dispatched as though they were marginal on an economic portion of their bid curve; however, the LMP prices at their locations were several dollars different than their offer price in some instances. These resources had the same schedule in as each other in these hours.[33] After a software patch was implemented, case 7a-3 was run and analyzed to confirm that both problems were corrected. We verified that in case 7a-3 there were no inconsistencies between bids, prices and schedules for any intertie resources.

In rerun IFM case 7a, the prices for one resource were inconsistent between two data files. The prices used in the dispatch do not appear to be consistent with the prices that are re-calculated using the underlying shift factor, shadow price, and penalty factor data. This kind of anomaly was not identified in rerun IFM case 110, but also occurred in case 7a-3. CAISO determined that this issue is related to a mapping issue within the master file and has been fixed in the master file; however, case 7a-3 was a re-run of an old case using old master file data.

There were several new issues found in the cases processed after the completion of the Interim Report. These issues are discussed below.

In case 110-3, we observed 1 resource that was economic for energy and was ramping up, but only ramped at 20-minutes, rather than 60-minutes. This unit was not providing reserves and ramped up more than 20-minutes in the next interval. The CAISO opened a variance on this issue. This issue was also seen in case 110-4. Further review of this anomaly has been postponed pending possible changes to the treatment of 20- and 60-minute ramp constraints.

In the wheeling1 case and the marketsim case, we observed a resource that was scheduled to its upper limit for energy when it is uneconomic by $0.06 and $0.18/MWh, respectively, for energy in each case. It is unclear why the unit did not ramp down in this interval. The CAISO opened a variance on this issue. It was conjectured that this was due to the case timing out before reaching an optimal solution. This issue was not seen in the wheeling2 case which was re-run with a longer time-out time.

[33] These anomalies were not present in rerun IFM case 110.

In the forbidden region case, we initially observed several resources that were ramping further than their ramp rates and forbidden region crossing time would allow. After the CAISO reviewed the methodology used for changing bid ramp rates and corrections were made, the issue was not seen in the forbidden2 case or forbidden3 case.

In the forbidden region case, we initially observed numerous instances of resources that were scheduled for regulation even though their total schedules were not inside their bid in regulation ranges. The CAISO opened a variance on this issue and the problem corrected. This issue was not seen in the forbidden2 case or forbidden3 case.

In the wheeling1 case, LECG observed one price capped load bid that was scheduled economically to a break point in its bid curve in the scheduling pass, but its pricing pass schedule was 0.1 MW higher than the scheduling pass schedule even though the margin for this 0.1 MWs was -$63.75. It is unclear why the load was scheduled higher in the pricing pass. It was conjectured that this was due to the case timing out before reaching an optimal solution. This issue was not seen in the wheeling2 case which was re-run with a longer time-out time.

In the marketsim case, an export bid was scheduled partially in the scheduling pass with an LMP at its location equal to the self-schedule penalty. In the pricing pass, the price was lower than the pricing pass penalty, but the export's schedule was not increased from its scheduling pass schedule. This apparent dispatch inconsistency was a result of the initial incorrect calculation of prices on export constrained interties. The dispatch was consistent with the correctly calculated price.

In case 110-3, the marketsim case, and the forbidden region case we observed resources that were scheduled to a level that was greater than their daily energy limit. Siemens explained that this was due to the penalty for violating energy limits being set to zero in this case. With the penalty set to correct value, no violations of daily energy limits were identified in the forbidden2 case or forbidden3 case.

A special case – fc_ifm_test_block_bid_e062908_s517268343 – was run to test the scheduling of block schedules in IFM. We observed that the block schedule was marginal over its block period in the scheduling pass, however, in the pricing run the average pricing pass price was greater than its bid over the block period. The resource's pricing run schedule remained equal to the scheduling pass schedule even though it appeared economic to be scheduled to a higher point. Siemens explained that for block transactions, the software initially blocked the pricing run schedule at the level of the scheduling run schedule regardless of the price in the pricing run. CAISO determined that it did not intend for block transactions to be fixed at the scheduling run level in the pricing run. A software change was made to correct this and the block2 case was run on the new software patch. We confirmed that block transactions are now correctly scheduled and are flexible in the scheduling and pricing runs.

In the forbidden2 case, we initially identified instances in which block transactions had offer prices that changed during a block period. This was a result of an error in the treatment of block transactions with one hour min run times. This was corrected and operated as intended in the forbidden3 case.

Our preliminary and interim reports noted that we observed a number of instances in which the scheduling of resources was not completely optimal but within the tolerance specified for the software (the MIP gap). While these instances continue to exist, changes that have subsequently been made in the way the objective function is defined have reduced the frequency of these instances as discussed below.

We initially observed a number of instances in the base cases and cases 7a, 102, 104, 106, 107, 108, 109, 110, and 111 in which units were dispatched to a breakpoint on their ramp capability curve, rather than to the point at which price equaled incremental cost or at which the unit would have been ramp constrained.[34] Thus, it would have been profitable to dispatch these resources to a different location on their energy bid curve.[35] These outcomes appear to be the result of the specified MIP gap, given the design of the software and the shape of the ramp rate curve submitted by these units. As a result of the complexities in the ramp rate curves, these units are treated as ramp constrained in the optimization although they actually should not be ramp-constrained at those dispatch points. In some instances, units are treated as ramp constrained in the dispatch without any change in dispatch from hour to hour.[36] All of the units treated as ramp constrained in this manner submitted ramp rate curves that can be described as W shaped, many associated with combined cycle modeling. These instances continued to be present in rerun IFM cases 7a and 110 but with reduced frequency.

There were 55 instances of MIP Gap in the wheeling1 case. These were associated with this case timing out before it could get to the .5% MIP Gap. The case

---

[34] .5% mip gap base case – 9 instances; high load base case – 13 instances; Case 7a – 4 instances; case 102 – 10 instances; case 104 – 6 instances; case 106 – 31 instances; case 107 – 3 instances; case 108 – 62 instances; case 109 – 8 instances; case 110 – 6 instances; case 111 – 20 instances. The number of MIP gap issues fell from 4 to 2 in the rerun of case 7a and from 6 to 1 in the rerun of case 110. Wheeling1 case – 55 instances; wheeling2 case – 14 instances; case 110-3 – 4 instances; marketsim case – 0 instances; forbidden region case – 8 instances; case 7a-3 – 0 instances; block2 case – 1 instance.

[35] This both includes instances in which units ramp to a break point and stop when it would have been optimal to ramp further and instances in which it would have been optimal to not ramp as far in that interval.

[36] At present, none of these units was flagged as ramp constrained in the output data, indicating that the ramp constrained flag is still not completely reliable.

was re-run with a longer time out period and in wheeling2 case there were only 14 MIP Gap instances.

In the .5% MIP gap base case there are two instances of MIP gap issues involving the scheduling of units for regulation. The units in question have different ramp rates depending on whether they are providing regulation and their commitment to provide regulation causes them to forgo larger profits in the energy market by reducing their ramp rate. One of these instances remained in the .01% MIP gap base case; the other resource was correctly scheduled in the case run at the lower MIP gap.

The instances of non-optimal dispatch identified above have been attributed to "MIP gap." The underlying issue is that when there are integer choices in the unit commitment and dispatch optimization problem, there is an inherent potential, given the resulting non-convexities, for the optimization to select a solution which is optimal given the choice of these zero one variables, but is not globally optimal. In the Siemens day-ahead market and RTUC software there are a number of such integer choices, involving unit commitment state, the ability of the resource to provide particular ancillary services, and the unit's ramp range. As a result, as noted above, there are a number of instances in which the solution is not globally optimal, and instances in which the dispatch is not optimal given the unit commitment. Virtually all of the MIP gap issues identified in this report arise, directly or indirectly, from the degree of flexibility in specifying ramp rates provided by the CAISO market design, and in many cases arises from the shape of the ramp rate curve specified by the market participant.

Our review has verified that these dispatch inconsistencies do not reflect erroneous price calculations, but simply reflect limitations on the optimality of the dispatch given the trade-off between performance and additional iterations. The prices are correctly calculated given the unit commitment, the ramp rate used in the dispatch solution, and the constraints which were binding in the dispatch solution. The calculation of LMP prices from a dispatch which is not fully optimal has precedent in PJM's operation from 1998 into 2002, when PJM had a limited set of dispatch tools and PJM calculated prices based on the dispatch, but the settlement prices were not everywhere consistent with the dispatch because the dispatch was not fully optimal.

The CAISO testing included runs of test cases using both 0.5% and 0.01% MIP gap. Review of the differences between these cases and the formulation of the objective function in the IFM software engine since preparation of our preliminary report identified elements of the calculation that were inflating the value of the objective function. The inflated value of the objective function interacted with performance criteria defined as a percentage of the value of the objective function (the "MIP gap") to cause iteration to stop further from the optimum that was intended by the design specifications. The formulation of the objective function was modified by Siemens, resulting in objective functions that typically have a much lower absolute value, implying a lower absolute mip gap for a given percentage standard. IFM cases 7a, 110 and the base case were rerun using the revised objective function formulation. The MIP gap in the reruns was much lower than in the original .5% MIP gap solutions in IFM cases 110 and the base case. This was not the case for the rerun of case 7a, but as noted previously, the case used to retest case 7a is not exactly the same case as the original case 7a, having different offer prices for some units.

RTUC

As in IFM, our review of the CAISO RTUC test cases has found that resources are dispatched consistently with their bids, offers and the LMP prices at their location in almost every instance, and none of the unresolved inconsistencies indicates the existence of a flaw in the calculation of prices. The inconsistencies that were identified and their resolution are described below.

In the basecase, we initially observed that a unit ramped up to 246 MW when its bid was more than $0.01 above the LMP price at its location, which is outside the rounding tolerance. We observed one import whose schedule was set in the scheduling pass and treated as though it were marginal, but its offer differed from the LMP by more than $0.01, which is again outside the rounding tolerance. Since there are small errors in replicating LMP prices in this case, both of these dispatch inconsistencies may reflect the impact of those price calculation errors. Similar occurrences of the dispatch and LMP being off by between $0.01 and $0.02, in cases in which we could replicate the calculated LMPs, occurred in cases 208(1) and 208(2). This issue was not observed in the rerun of the RTUC basecase or in cases 203 and 204.

There continue to be resources that appear to be dispatched uneconomically, and that we can manually calculate to be ramp constrained, but that are not listed as ramp constrained in the emm_scuc_output_valid_ramp.csv file. This occurred in the base case and in all the additional RTUC cases in IFM. Use of the cmm_scuc_output_ bid table to identify ramp or limit constrained units has resolved some but not all of these instances. This issue was present in rerun RTUC cases 202 and 203. These are not pricing errors but these kinds of omissions will slow price validation once MRTU is implemented.

We observed 10 units that were dispatched to breakpoints in their ramp rate curve, rather than to their economic dispatch point, in the .5% MIP gap basecase and 4 such units in the .01% MIP gap basecase.[37] Similar instances were found in all the other RTUC cases. As discussed above in the context of IFM, these reflect the effect of the

[37] There were seven instances in the rerun of the base case (case 202).

non-convexities in these ramp rate curves and the way the objective function was originally defined. A related instance occurred in case 204, in which a unit with a regulation schedule ramped to a breakpoint in its energy dispatch ramp curve, while it would have been more profitable to continue ramping down. While units ramping to a breakpoint on their ramp capability curves is a normal MIP gap outcome, in this case the resource was stopping at a breakpoint on the wrong ramp capability curve. These MIP gap issues were also observed in the rerun of cases 203 and 204.[38]

We also observed in cases 201, 204, 207, 209 and 210 units with self-schedules for energy that were not dispatched consistent with their self-schedules, although the price in both the pricing and scheduling passes exceeded the self-schedule penalty price. In case 210, the minimum and maximum energy levels were missing from the bid data, causing the energy self-schedules to be reset to 0 MW. This issue was not observed in rerun RTUC cases 202, 203, or 204. We also observed units in case 210 that did not submit an energy self schedule, but received one nonetheless. This was because the units had Must Run designations for regulation and the software will create an energy self-schedule at the units minimum regulation range for such units.

We observed an instance of a unit with a $0 incremental energy bid and a $0 no load cost being decommitted in case 207.

We noted that one unit in case 209 was constrained by its regulation ramp rate even though the unit was not providing regulation and would have been more profitable had it been scheduled using its energy ramp rate. This unit is listed as being on regulation in EMS so it is subject to the more constraining of it regulation and energy ramp rates during the current hour. Siemens has confirmed that the software is implemented such that EMS regulation status affects both hours spanned by an RTUC interval, while the CAISO has stated that this status should only affect the ramp rate during the current hour. There is on-going discussion of this issue. Instances of EMS regulation status affecting ramp rates during the second hour of an RTUC interval were also observed in rerun RTUC cases 202 and 203 and in case 206b.

We observed three instances in Case 210 (a HASP case) of imports, which should have been scheduled to the same level in all four intervals of an hour, being uneconomically scheduled in the first interval of the hour. The CAISO has submitted a variance on this issue.

In case 206b, four pre-dispatch imports were scheduled down into their self schedules during the scheduling pass for the second hour of the case. However, in the

---

[38] The number of such MIP gap issues was unchanged in the rerun of case 203 and rose from 8 to 9 in the rerun of case 204.

pricing pass, their schedules returned to the full amount of the self schedule for two non-consecutive intervals though the schedule change was uneconomic, and also caused the schedules for the pre-dispatch imports to vary over the hour. The CAISO was to create a variance on this issue.

There was one load schedule anomaly in Case 210. An export was scheduled at 0.1 MW when the price in the scheduling run was above the export's bid. The CAISO has written a variance on this issue.

In the export case, there were instances in which where imports have a 0MW schedule even though they had an operating mode of must-run, were listed as online, and had a non-zero self-schedule. In all cases, the LMP at these resources' locations were greater than the -$30 self-schedule penalty in the pricing pass. This issue was resolved with patch 20065, and tested in the stuc and u2_lap cases.

In the export case, there were also instances in which exports have a 0MW schedule even though they had an operating mode of must-run, were listed as online, and had a non-zero self-schedule. In all cases, the LMP at these resources' locations were less than the $500 self-schedule penalty in the pricing pass.[39] This issue was also resolved with patch 20065 and tested in the stuc and u2_lap cases.

There were 48 instances in the export2 case in which resources were scheduled to turn off before the end of their minimum run time. In all cases these resources had an initial status of online and turned off in the first interval of the RTUC run even though the minimum run time had not been satisfied. The way minimum run time constraints are honored across RTUC intervals is being reviewed by the CAISO and Siemens.

We observe an inconsistency in the resource limits for RTUC that are output from the software. In the lap3 case these limits appear to constrain the energy plus upward ancillary service schedules, however in the export2 case this limit constrains the energy schedule and spin is scheduled above the limit.

There were wheeling transactions present in the export case and export2 case. In all instances, valid wheeling transactions were scheduled correctly. These cases contained some invalid wheeling transactions, unbalanced wheels, because the data used to create the case was not entered through SIBR.

---

[39] A perhaps related issue was observed in the class B case 2006-4 discussed in section III. In this case, there is one export that is economic to be scheduled higher, but is not. The bid price for the export exceeds the price at that location by $410, so the export is clearly economic. There do not appear to be any unit derates on the export load and the export load is not at a limit in the output bid file.

RTD

We identified a number of minor issues involving the scheduling of units in RTD but none of these issues indicated the existence of a flaw in the calculation of prices. The issues that have been identified are described below.

We observed 12 instances in the base case of units providing regulation in real-time, but not having a regulation schedule, ramping at the lower of their energy ramp rate or their regulation ramp rate as if considered to be providing regulation in intervals in the next hour in which they also lack a regulation schedule. CAISO has reviewed the requirements and determined that this is appropriate.

We noted two instances in the base case of a unit's energy plus ancillary services up schedules being greater than its maximum capacity. We observed 10 instances of this in case 301, four instances in case 303 and 22 instances in case 304. The CAISO has an outstanding variance on this issue.

We initially identified a unit in case 301 that is awarded a 1 MW schedule for regulation down in the RTUC market, but is scheduled for energy at 0 MW in RTD. The regulation down schedule was therefore disqualified. The same pattern occurred on one unit in case 304. CAISO identified this as a data issue, reran case 301, and the unit received a 2 MW energy schedule and a 1 MW regulation down schedule.

Four units were dispatched in the scheduling pass of case 301 in a manner inconsistent with the scheduling pass prices. One similar instance was identified in case 303. All of these units had self-schedules that were modified in the scheduling pass. There do not appear to be any pricing inconsistencies in the pricing pass, so the settlement prices were correctly calculated given the constraints on modification of the self-schedules in the pricing pass.

We identified one unit in case 303 that was ramping down over the entire time horizon, although its operation at a higher output level was economic in every interval and the unit was not constrained by a changing maximum limit. The CAISO submitted a variance on this issue.

There were no dispatchable loads and hence no load schedule anomalies in any of the RTD cases.

We identified a data issue all the RTD cases in which units are ramp constrained, but are not flagged in the data as ramp constrained. This occurred 24 times in the base case, three times in case 301, and three times in case 304. This issue does not reflect any price calculation error or dispatch issue but should be corrected prior to market implementation to speed price verification.

We observed six instances attributable to “mip gap” in the base case in which units were dispatched to a breakpoint on their ramp capability curve, rather than to the point at which price equaled incremental cost or at which the unit would have been ramp constrained. There were five occurrences of this in case 301 and four in case 304.

Conclusion

While we have identified a number of inconsistencies between the calculated prices and the dispatch in various IFM, RTUC and RTD test cases, none of the inconsistencies reveal any kind of fundamental error in the calculation of prices in these cases. All of the issues identified appear to involve inconsistent truncation or rounding or minor imperfections in the dispatch, not fundamental pricing issues.

4. Replicate Ancillary Service Prices from Shadow Prices

IFM

As a result of reporting changes, a manual process was used to replicate the ancillary service prices for the .5% MIP gap base case, the high load base case, cases 7a, 102, 104, 106, 107, and 108. This manual calculation derived prices for regulation up by subtracting the SAS generated spinning reserve price from the regulation up price, and arrived at a spinning reserve price by subtracting the non-spinning price from the spinning reserve price. We then compared these calculated prices with the prices presented in the CAISO's emm_scuc_output_bid.csv file. No errors were identified when we performed this analysis on the base cases, case 7a, 102, 104, 106, 107, or 108. We accounted for the reporting charges and incorporated this test into the SAS price validation tool and it identified no errors for cases 109, 110, 111, or rerun cases 7a and 110. There were no errors in the wheeling1 case, wheeling2 case, marketsim case, forbidden region case, forbidden2 case, or forbidden3 case.

An issue regarding the reporting of incorrect $0 ancillary service clearing prices in the emm_scuc_output_bid.csv file for units that were not scheduled to provide ancillary services was resolved prior to the preliminary report. In cases 7a, 102, 104, 107, 108, 109, 110 and 111, this file is populated with the correct clearing prices for all commodity types that were offered by a resource. However, in case 110 and rerun case 110, LECG was unable to verify that ancillary service prices were correctly calculated because no shadow prices were reported for ancillary service regions in which no resource offered ancillary services.

IFM Case 110 was designed such that there are shortages of ancillary services at any price both in the CAISO as a whole, and in some subregions. In the initial run of case 110 in some hours, ancillary service prices were set by shortage values and shortage values were cascading from region to region,[40] but in other hours in which the ancillary service could not be met, this cascading did not appear to be implemented and prices were set by the highest offer price rather than shortage values.

The apparent cause of these anomalies was that the version of the IFM and RTUC software used to run case 110 carried out an initial calculation to identify instances in which insufficient ancillary services are offered to meet a regional requirement and then relaxed the requirement to be equal to the amount offered. This accounts for the instances in which there was a shortage of ancillary services within a region but the price was set by the highest offer price, rather than by the shortage value. Prices were sometimes set by the shortage values, however, because not all shortages of ancillary services were identified in this initial calculation. The CAISO requested that the vendor remove this feature and achieve the intended result by setting the penalty value for ancillary service shortages to zero in the pricing pass. Several test cases were run to test whether this change in the handling of ancillary service shortages was correctly implemented. Test case 110-4 verified that a code change produced the intended result in a case with cascading turned off across ancillary service products, but the code change has not yet been implemented in a software patch.[41]

In verifying that the cascading of ancillary service prices was working as intended, we observed that on-line units are able to submit offers to provide both spinning reserve and non-spinning reserve from the same capacity at distinct offer prices. Moreover, the entire rampable capacity of such a unit could be offered either to provide spinning reserve or non-spinning reserve with the overall ramp limit enforced when the market was cleared. Price cascading therefore operates somewhat differently than in other markets, as rampable capacity on an on-line unit could clear either as spinning or non-spinning reserve and would only trigger cascading if cleared as spinning reserve. Price cascading operated correctly in the base cases, case 7a, 102, 104 and 108, given these features. We observed instances in the test case bid data in which non-spinning reserves were offered at higher offer prices than spinning reserves on the same unit and

---

[40] For example, in case 110 there is insufficient spinning reserve scheduled to satisfy the CAISO spinning reserve requirement in all 24 hours and there is also insufficient spinning reserve scheduled in regions 2 through 5 in all 24 hours. With cascading based on shortage values, the price of spinning reserve for region 2 would be the sum of the shortage value for spinning reserve for the CAISO region and for region 2. This cascading is observed in the scheduling and pricing passes in hours beginning 8 through 10 and 14 through 18 PST, in which the shadow price of spinning reserves is $12,000 for the CAISO region in the scheduling pass and $1000 in the pricing pass. However, in the remaining hours, it appears that the shadow price of spinning reserve in the CAISO region is set by the highest accepted offer price.

[41] The export case also has violated ancillary service requirements and has some dispatch inconsistencies that likely reflect incorrect calculation of ancillary service prices that the CAISO will want to reexamine once a software patch is available.

instances in which non-spinning reserves were offered at lower offer prices than spinning reserves. There were no instances in the test cases in which non-spinning reserve prices exceeded spinning reserve prices, however, it may be possible for market participant offers to produce this outcome.

RTUC

We were able to recalculate all ancillary service prices in the RTUC base case and in cases 210 and 206b, however, all spinning and non-spinning reserve shadow prices were equal to zero. We were also able to recalculate all ancillary service prices in cases 201, 203, 204, 205, 206, 207, 208(1) and 208(2) and in the rerun of the RTUC basecase and cases 203 and 204. The spinning and/or non-spinning reserve shadow prices were not equal to zero in these cases.

RTD

The LECG SAS price validation tool does not test the recalculation of ancillary service prices based on shadow prices in RTD because RTD does not schedule ancillary services.

5. Validate Ancillary Service Prices based on Marginal Offers

IFM

We were able to identify the correct number of marginal ancillary service suppliers in IFM case 106. A marginal ancillary service resource was identified for every binding ancillary service constraint in the rerun of case 7a. This analysis was not carried out for the rerun of case 110 because of the inconsistencies involving the way ancillary service shortage values appear to be setting price as noted in the preceding section. The correct number of marginal resources was identified in case 110-4.

RTUC

We are able to identify a marginal ancillary service resource for every binding ancillary service constraint in the reruns of RTUC cases 202, 203 and 204.

In case 208(2), the regulation down requirement was not being met and the price of regulation down in the pricing pass was $250, far above the highest accepted bid of $20. In both 208(1) and 208(2), the regulation up requirement was not met, and the regulation price was set by the highest accepted bid. As in the case 110 rerun, these inconsistencies probably result from the way the ancillary service requirements were relaxed when the requirement cannot be met and the pricing run penalty is set to zero.

RTD

The LECG SAS price validation tool does not identify marginal ancillary service suppliers in RTD because RTD does not schedule ancillary services.

6. Examine Ancillary Service Prices for Consistency with Ancillary Service Schedules, Energy Schedules and Energy Prices

IFM

Our review of the base cases and analysis track cases determined that resources scheduled to provide ancillary services are being scheduled to provide an amount that is consistent with the ancillary service prices, and energy market opportunity costs, in almost every instance. None of the inconsistencies identified revealed a flaw in the calculation of ancillary service prices. The inconsistencies that were initially identified and their resolution are noted below.

We initially observed seven instances in case 108 in which the co-optimization of energy and ancillary services appeared to be incorrect by more than \$0.01, but less than \$0.02, based on the reported energy and ancillary service prices. These anomalies were likely related to the large constraint shadow prices in case 108, combined with the rounding of transmission constraint shift factors discussed above in Section 3. This also occurs in three instances in case 109 and three instances in case 110. After Siemens corrected the inconsistent rounding, this issue was not observed in rerun cases 7a or 110.

We initially observed instances in cases 109, 110, and 111 in which offline units were scheduled to provide amounts of non-spinning reserves that were not feasible given the resource's start-up time, lower operating limits, and non-spin ramp rates.[42] This issue was corrected and was not observed in rerun IFM cases 7a or 110.

We observed instances in case 108 in which units were not scheduled for their full self schedule for non-spinning reserves, because they had an economic bid with a larger margin than the self-schedule modification penalty for non-spinning reserves. We believe that the software is scheduling these units correctly based on economics.

In the rerun IFM case 7a, we observed instances in which units were scheduled non-optimally between two ancillary services. These units were scheduled for one ancillary service up to the resources 10-minute ramp limit and then scheduled to provide

---

[42] It is our understanding that offline resources cannot be scheduled to more than their: lower operating limit + (10 – start-up time)*non-spinning ramp rate. This issue occurred in 27 instances in case 109, 19 instances in case 110, and 93 instances in case 111.

another lower valued ancillary service for a fraction of a minute of ramp capability. We did not see this issue in rerun IFM case 110.

In the class B case 2006-4, we have observed some resources that are not providing regulation, but whose schedules are constrained by the upper regulating range, although they is economic to be scheduled to a higher energy output based on their offer prices. These resources are not at any other limits.

In the forbidden region case, we observed a new anomaly involving numerous instances of resources that were scheduled for regulation even though their total schedules were not inside their as-bid regulation ranges. The CAISO opened a variance on this issue and it was corrected. This issue was not seen in the forbidden2 case or forbidden3 case.

We observed one resource in the block2 case that was scheduled for non-spin during its minimum down time. This unit did not have a self-schedule for non-spin. CAISO opened a variance on this error and fixed it but we have not yet received a rerun.

While resources scheduled to provide ancillary services are almost always scheduled to provide the correct amount of ancillary services, we have identified instances in which resources were not scheduled to provide ancillary services when it would have been economic for them to have done so. These MIP Gap issues in the base case, and cases 7a, 102, 104, 106, 107, 108, 110, and 111 arise from the binary commitment decision for regulation (commitment = 0 or 1), as well as the commitment decision for non-spinning reserves when the unit is offline (commitment = 0 or 1). These binary decisions can cause units not to be scheduled for regulation or non-spinning reserves when they are in fact economic to provide that ancillary service. In the base cases, and cases 7a, 102, 104 and case 108, a varying number of units were scheduled to provide 0 MWs of an ancillary service, were not committed to provide that ancillary service, but it would have been profitable for the units to have provided that ancillary service.[43] These generally appear to be instances of non-optimal commitment due to the MIP gap.

One of these instances in case 108 included a resource that was forgoing a margin of nearly $1000/MWh by not being scheduled to provide regulation. This rather large departure from the optimal dispatch appears to have been a result of the value of the objective function for this case. The CAISO reran this case at a .01% MIP Gap level to

---

[43] 9 instances in the base case, 92 in the high demand base case, 7 in case 7a, 25 units in case 102, 7 in case 104, 38 in case 106, 29 in case 107, 113 in case 108, 2 in case 110, and 18 in case 111, 2 in rerun case 110, 5 in rerun case 7a, 74 in wheeling1 case, 13 in wheeling2 case, 16 in marketsim case, 23 in forbidden region case, 1 in case 110-3, 8 in case 7a-3, and 27 in the block2 case.

determine if these instances of non-optimal scheduling would disappear. The number of instances of non-optimal ancillary service scheduling in case 108 fell from 113 to 28 at the lower MIP gap level, and the instance of the unit forgoing the large regulation margin that was present in the .5% MIP gap case was eliminated with the unit correctly scheduled to provide regulation. The changes that were made to the way the objective function is defined addressed the potential for such material departures from optimality by materially reducing the absolute MIP gap.

There were 74 MIP gap-related dispatch errors in the wheeling1 case. These were associated with this case timing out before it could get to the .5% MIP Gap. The case was re-run with a longer time out period and in wheeling2 case there were only 13 MIP Gap instances.

RTUC

We observed in the RTUC base case and in case 210 that there continued to be instances of resources that are in their minimum down time being scheduled to provide non-spinning reserves. This situation appears to arise because minimum down time constraints are relaxed on units that are self-scheduled to provide non-spin and these resources have 0 MW non-spin self-schedules. With the minimum downtime constraint relaxed, these units received non-zero non-spin schedules based on their economic bids. CAISO has determined that the minimum down time constraint should be relaxed in this situation and therefore, the software is acting as intended.

We noted one instance in case 203 of an off-line unit being scheduled to provide non-spinning reserves although it had a 30-minute startup time. This issue was not observed in the rerun of the RTUC basecase, nor in the rerun of cases 203 or 204.

We noted an issue in cases 203, 204 and 206 in which units were offering more megawatts of spinning reserve than they could provide. CAISO indicated that this was a result of how the test data were prepared.

We observed one unit in case 210 that was not awarded its self-schedule of regulation down, although the unit was in its regulating range, was economic, and had the capacity to provide it. The resource was qualified to provide its full self-schedule in the day-ahead market, but was only scheduled partially in the day-ahead market. In this case, it appears that the RTUC schedule was limited to the quantity that cleared in the day-ahead market, rather than the quantity that was it was qualified to provide. CAISO is investigating why the RTUC dispatch appears to be constrained to be consistent with the day-ahead market schedule in this instance. This was also observed in rerun RTUC cases 202 and 204.

LECG identified four instances in case 204 of offline units not being scheduled to provide non-spinning reserves, despite having a non-spinning reserves self-schedule and not being in their minimum down times. Such instances were not observed in the rerun cases 202, 203, and 204.

In cases 201, 204 and 207 and 210, we identified instances of units scheduled to provide regulation despite being scheduled to generate energy outside their regulation ranges. CAISO has determined that there is a hierarchy of choices the software can make for regulation status and all of the apparent anomalies in these and other cases are consistent with the intended rules.

In rerun RTUC case 204, we observed instances in which units were scheduled non-optimally between two ancillary services. These units were scheduled to provide one ancillary service up to the resources 10-minute ramp limit and then scheduled to provide another lower valued ancillary service using a fraction of a minute of ramp. This issue was also seen in the rerun IFM cases and Siemens explained that it is related to a small threshold added to the ramp time of the combined upward AS ramp time allowed. CAISO and Siemens believe that this threshold should be removed, but are verifying that there is not a reason for it before doing so.

As in IFM, we observed instances of non-optimal commitment of resources providing ancillary services. There were four such MIP gap observations in the RTUC base case.[44] There was one example of this MIP Gap issue in case 203, 62 examples in case 204, two examples in case 207, three examples in case 209 and two examples in case 210. There were no such instances in the rerun of RTUC basecase or in cases 203 and 204.[45]

RTD

The LECG SAS price validation tool does not review ancillary service prices for consistency with ancillary service schedules in RTD because RTD does not schedule ancillary services.

---

[44] The number of such instances fell from 4 in the .5% MIP gap base case to 2 in the .01% MIP gap base case.

[45] Thus, the number of MIP gap non-optimalities fell from 62 in the original 204, to zero in the rerun with the new objective function.

7. Examine Unit Commitment for Consistency with LMP Prices

IFM

The calculation of uplift costs by resource was not carried out for the November cases, but was included in the final evaluation of analysis track cases. The uplift evaluation was carried out considering both pricing pass and scheduling pass prices.

There were no instances of substantial uplift costs on units committed on economics in either of the base cases, cases 7a, 102, 104, 108, 109, 110, or 111. There were no substantial uplift costs in the wheeling1 case, wheeling2 case, marketsim case, or forbidden region case. There were instances of self-scheduled units being committed uneconomically based on their pricing pass penalty values, but all of these units were correctly committed based on their self-schedules in the scheduling pass. Uplift charges on other units were within the range that we would normally expect to see in such a unit commitment problem.

There were substantial uplift costs on 4 resources in case 110-3, over $200,000 per resource. These resources were not must-run, self-scheduled, nor were they providing a large amount of reserves (and in two cases no reserves). CAISO provided a re-run of this case and we confirmed that overall production cost increased when one of the 4 resources was de-committed for the entire horizon, indicating that the original solution was lower cost, despite the large uplift.

We calculated large forgone energy revenues for one resource in case 107 and 7 resources in case 106. All of these resources are among the resources for which we were unable to replicate the calculated LMP prices and all are indicated in a message file to be electrically disconnected from the grid. In some instances, the resources were scheduled to provide non-spinning reserves despite not being connected to the grid. As noted above, the CAISO and Siemens are reviewing various aspects of the modeling of these resources.

In the wheeling1 case, large forgone energy revenues were calculated for many resources. It is believed that these were associated with this case timing out before it could get to an optimal dispatch. The case was re-run with a longer time out period and in wheeling2 case the level forgone energy revenues were in line with expectations

The rerun of the IFM base case and 110 using an improved formulation of the objective function resulted in lower absolute MIP gaps than in the original cases. The absolute value of the MIP gap in the base case rerun was less than half the value in the .01% MIP gap case and far lower than in the .5% MIP gap case. The prices in the basecase rerun were similar to those in the .01% MIP gap case and the prices for one

LAP were materially lower than in the original .5% MIP gap case and were consistent with those in the .01% MIP gap case.

We calculated relatively large forgone energy revenues for two resources in rerun IFM case 7a. CAISO has indicated that these resources are located in a small load pocket, such that the LMP prices within the pocket change drastically if the units are online or offline. The resources at issue were committed during some hours in case 7a and their operation was slightly uneconomic during those hours.

RTUC

There was one unit with a must run operating mode that was not committed for energy in any of the RTUC test cases (except case 206b), including the reruns. This appears to have been the result of a data issue in which the unit was not modeled as connected to the grid in any of these cases.

In the export case, export2 case, and lap3 case there were 50, 249, and 38 instances, respectively of self-scheduled resources that were not committed and had a status of "offline." All of these resources were in their minimum down time. It is our understanding that self-schedule resources should always be committed even if in their minimum down time.

RTD

The SAS price validation tool does not review unit commitment in RTD because RTD is not commitment software.

III. CLASS B

The CAISO also asked us to carry out a limited review of several class B cases. These cases were structured to test the performance of certain elements of the IFM or RTUC software under particular conditions. For these cases, we did not carry out the complete set of tests described in Section IIA; instead we:

1. Verified that the CAISO's test methodology was conceptually appropriate for the test objectives.

2. Verified the CAISO's observations regarding the test results using the test case data provided by the CAISO or, if necessary, by obtaining additional data.

Cases 2001 and 2002 were designed to test whether the relative constraint relaxation priorities between specific hard constraints (unit minimums) and specific penalty priced constraints (transmission limits and self-schedules above the unit minimum) were operating as intended such that the hard constraints would be enforced in the scheduling pass while the penalty constraints would be relaxed. The software operated with the correct priorities in both cases, enforcing the unit minimums, curtailing self-schedules, and violating the transmission limit.

Cases 2003 and 2004 tested whether the priority of transmission constraint relaxation was correctly applied between branch group constraints and line constraints with different penalty prices. The correct priority for relaxation of constraints was observed in the scheduling pass for constraints with different penalty prices.

Case 2005 tested whether if no self-schedules are curtailed in the scheduling pass and no other constraints are violated (e.g., there no transmission constraints are relaxed), overall energy and ancillary service schedules will be the same between the pricing pass and the scheduling pass. This outcome was observed. As expected, in instances in which there were multiple schedules with identical bid or offer prices at the same location, the specific schedule accepted could change between the scheduling and pricing passes.

Case 2006 tested whether LAP clearing prices are correctly set when there are uneconomic adjustments and that the LAP price can be recalculated from the underlying constraint shadow prices, shift factors and loss penalty factors. This could not be confirmed until the LAP price calculation issues were resolved. Now that they have been resolved, Case 2006 was re-run[46] and the Pnode and Apnode recalculation analyzed. We identified the Pnode issue associated with the incorrect sign being applied to the export constrained shadow price. There was a Pnode mapping issue relating to a Pnode that was mapped to two interties. We found one Apnode recalculation error related to the above Pnode export error. We identified an issue where it appears small Apnode participation factors are being dropped from the analysis and treated as disconnected nodes. CAISO and Siemens were looking into this issue.

The CAISO provided a third case 2006[47] that was run on a software patch with the fix for the incorrect sign getting used on the export constrained shadow prices. We confirmed that the pricing and dispatch of export constrained nodes was working correctly. The third case 2006 still contains the Pnode mapping issue where two resources on the same Pnode are not linked to the same inter-tie location. Also, we continue to identify an issue where it appears that small Apnode participation factors are being dropped from the analysis and treated as disconnected nodes. CAISO and Siemens were looking into this issue.

---

46 fc_ifm_test_case2006_lapadj_p353_e081508_517270244

47 fc_ifm_test_case2006_p200020_e090108_517280540

A fourth 2006 case[48] was provided on a new patch. We confirmed that small Apnode participation factors are no longer dropped from the Apnode price recalculation. We identified one price recalculation issue on a Pnode where reported prices differed by $30 from the calculated price. We identified a large number of dispatch inconsistencies in the pricing pass involving self-scheduled load bids which were curtailed in the scheduling pass. These LAP load bids were dispatched consistent with the intended prorata curtailment in the scheduling pass which produced the relevant day-ahead schedules. However, some load bids were dispatched materially above their scheduling pass level in the pricing pass, although the pricing pass LMP at their LAP exceeded (by more than $1000 per megawatt in some hours) the $500 pricing run bid cap. This apparently relates to the way prorata curtailment was implemented in the pricing pass. Analysis of this case is hindered by an extremely large MIP gap, greater than 80%, and units which apparently should have been committed but were not. It has not yet been determined if there is some error in setting up the case that is leading to these problems.

Conversely, in the 2006-4 case, there are load resources at the same lap whose self-schedules were correctly curtailed in the scheduling pass, but then were dispatched materially down (i.e. more than just an epsilon) below their scheduling pass level in the pricing pass. These loads should not have been dispatched more than one epsilon below their scheduling pass level in the pricing pass. This observation also appears to be related the way the pro-rata curtailment was implemented in the pricing run.

Case 2012 validates LMP prices on both sides of a transmission constraint based on shift factors and shadow prices. This outcome was observed.

Cases 4004 (DLS) and 4004 (DGS) were used to compare the LMP decomposition between a Distributed Load Slack bus (DLS) and a Distributed Generation Slack bus (DGS). LECG only analyzed these cases from the standpoint of identifying price calculation errors. LECG was able to recalculate all of the Pnode prices. However, we were unable to recalculate some of the Apnode prices, including each of the default LAP prices in one or more hours. Now that the Apnode price recalculation issues have been resolved, it should be possible to rerun these cases and confirm the expected result. Case 4004 was re-run and we were able to recalculate Apnode prices, including all default LAP prices.

Cases 5003 and 5004 included changes in transmission limits within the time frame of the optimization. We verified that the change was correctly applied and reflected in the dispatch and prices.

---

48 fc_ifm_staging_case2006_e10132008_417210560