Experience with PJM Market Operation, System Design, and Implementation

Andrew L. Ott

Invited Paper

Abstract—This paper outlines the fundamental features of the PJM day-ahead energy market and real-time energy market. The Day-ahead market is based on a voluntary least-cost security constrained unit commitment and dispatch with several fundamental design features that ensure the market is robust and competitive. This market offers market participants the option to lock in energy and transportation charges at binding day-ahead prices. The flexibility of the day-ahead market rules provide all participants with equal access to the day-ahead market through consistent price signals and by providing all participants with the ability to submit virtual demand bids and virtual supply offers. These mechanisms promote liquidity in the markets. Economic incentives drive the convergence of the day-ahead and real-time market prices. The real-time energy market is based on security-constrained economic dispatch and is cleared based on the actual system operating conditions. The LMP-based markets support reliable grid operations through efficient price signals.

Index Terms—Electricity market, locational marginal pricing.

I. INTRODUCTION

PJM OPERATES the world's largest competitive wholesale electricity market and one of North America's largest power grids. PJM currently coordinates a pooled generating capacity of more than 67 000 MW and operates a wholesale electricity market with more than 200 market buyers, sellers and traders of electricity. The PJM market covers all or parts of PA, NJ, MD, DE, OH, VA, WV, and the District of Columbia. With the April 1, 2002, addition of PJM West, for the first time nationally two separate control areas now operate under a single energy market, single security-constrained economic dispatch and a single governance structure across multiple North American Electric Reliability Councils.

The PJM energy market consists of two markets—a day-ahead market and a real-time balancing market. The day-ahead market is a forward market in which hourly clearing prices are calculated for each hour of the next operating day based on generation offers, demand bids, virtual supply offers, virtual demand bids and bilateral transaction schedules submitted into the day-ahead market. The day-ahead energy market is a voluntary bid-based market that is cleared using a security-constrained unit commitment and economic dispatch.

The balancing market is the real-time energy market in which the clearing prices are calculated every 5 min based on the actual system operations security-constrained economic dispatch. Separate accounting settlements are performed for each market, the day-ahead market settlement is based on scheduled hourly quantities and on day-ahead hourly prices, the balancing settlement is based on hourly integrated quantity deviations from day-ahead scheduled quantities and on real-time prices integrated over the hour. The day-ahead price calculations and the balancing (real-time) price calculations are based on the concept of locational marginal pricing (LMP).

II. DAY-AHEAD ENERGY MARKET

The day-ahead market provides market participants with the ability to purchase and sell energy at binding day-ahead prices. It also allows transmission customers to schedule bilateral transactions at binding day-ahead congestion charges based on the differences in Locational Marginal Prices between the transaction source and sink locations. Load serving entities (LSEs) may submit hourly demand schedules, including any price sensitive demand, for the amount of demand that they wish to lock-in at day-ahead prices. Any generator that has entered into an installed capacity contract must submit an offer schedule into the day-ahead market even if it is self-scheduled or unavailable due to outage. Other generators have the option to offer into the day-ahead market or into the real-time market. Transmission customers may submit fixed, dispatchable or "up to" congestion bid bilateral transaction schedules into the day-ahead market and may specify whether they are willing to pay congestion charges or wish to be curtailed if congestion occurs in the real-time market. All spot purchases and sales in the day-ahead market are settled at the day-ahead prices.

After the day-ahead market bid period closes, PJM calculates the day-ahead schedule based on the bids, offers, and schedules submitted based on least-cost, security constrained unit commitment, and dispatch for each hour of the next operating day. The day-ahead market clearing process incorporates PJM reliability requirements and reserve obligations into the analysis. The resulting day-ahead hourly schedules and day-ahead LMPs represent binding financial commitments to the market participants. Financial transmission rights (FTRs) are settled at the day-ahead LMP values.

A. Market Design Objectives

As stated previously, the PJM market design is based on the concept of LMP. A key feature of an LMP model is that there is a fundamental consistency between the energy price and the price of delivery on the transmission system. In this model, the energy price difference between the injection point and the withdrawal point is equal to the transmission congestion cost.[1] Therefore, a market participant who injects (sells) energy at location A and withdraws (purchases) at location B will pay exactly the same as a market participant who pays the transmission congestion charge to deliver a bilateral contract from A to B. This consistency is a feature of both the PJM day-ahead market and the PJM real-time market.

In addition to the LMP concept, the fundamental design objectives of the PJM day-ahead energy market are: 1) to provide a mechanism in which all participants have the opportunity to lock in day-ahead financial schedules for energy and transmission; 2) to coordinate the day-ahead financial schedules with system reliability requirements; 3) to provide incentive for resources and demand to submit day-ahead schedules; and 4) to provide incentive for resources to follow real-time dispatch instructions.

The first market design objective is accomplished by providing a variety of alternatives for participation in the day-ahead market. The participation options include the ability to self-schedule resources, the ability to submit bilateral schedules and the ability to submit offers to sell or bids to buy from the day-ahead spot market. This flexibility ensures that all market participants have equal access to the day-ahead market. Therefore, any barriers to trade in the day-ahead market are minimized so that the market will be as competitive as possible. In order to further promote liquidity, the market design also includes the ability to submit purely financial positions in the form of virtual supply offers and virtual demand bids. In this way, the day-ahead market provides both the ability the hedge physical delivery and the ability to enter financial positions into the market. All positions that are cleared in the day-ahead market are financially binding and will liquidate in the balancing market if they are not covered by a real-time energy delivery.

The second market design objective is important to ensure that the day-ahead schedules are physically feasible and are consistent with reliable system operations. This feature is significant because it requires that the powerflow model used to analyze the day-ahead market is consistent with the powerflow model that is used in real-time system operations. It also requires that the Day-ahead market is cleared considering the same single contingency criteria and transmission equipment ratings that are used in real-time operations. Since the underlying powerflow model and operating constraints are consistent between the day-ahead forward market and the real-time dispatch, the LMP signals are consistent between the day-ahead and real-time markets as well. In addition to the powerflow model consistency, the day-ahead market also respects system reserve requirements and the generator

physical operating limitations.[2] This design feature ensures that the financial schedules that result from the day-ahead forward market are consistent with the physical transmission capability. Therefore the day-ahead scheduling process ensures that the transmission capability is not over-subscribed and ensures that the generation schedules are consistent with the generator's physical capabilities. The fundamental consistency between the forward market and the real-time market ensures a robust market design that promotes economic efficiency and it enables the market to avoid the gaming opportunities that have plagued other market designs.

The third market design objective involves more than just the fundamental structure of a two-settlement system.[3] It also requires that there is consistency between the market pricing mechanisms and that price convergence occurs between the markets over time.

The following series of examples will illustrate the economic incentives that exist in the two settlement design.

Example 1: A customer submits a day-ahead demand bid that clears for 100 MW at an LMP of U.S.$ 20/MWh. In the real-time market, the customer has 105 MW of demand and the LMP value is U.S.$ 23/MWh. In this case, the customer's resulting financial settlement is: day-ahead load payment = U.S.$ 20/MWh * 100 MW = U.S.$ 2000, Real-time balancing payment = (105–100) MWh * $23/MWh = U.S.$ 115. The total payment is U.S.$ 2115. If the customer had locked in his total load of 105 MWh at the Day-ahead price of U.S.$ 20 then his total payment would have been 105 MWh * $20 = $2100.[4] Therefore, in this case the customer has the incentive to submit the actual expected demand into the day-ahead market. This incentive is driven by the fact that the customer experienced a demand increase between the day-ahead and real-time markets at the same time that the entire market exhibited the same trend as indicated by the higher real-time LMP.

Example 2: A customer submits a day-ahead demand bid that clears for 100 MW at an LMP of U.S.$ 20/MWh. In the real-time market, the customer has 95 MW of demand and the LMP value is U.S.$ 23/MWh. In this case, the customer's resulting financial settlement is: day-ahead load payment = U.S.$ 20/MWh * 100 MW = U.S.$ 2000, Real-time balancing payment = (95–100) MWh * U.S.$ 23/MWh = −$115. The total payment is U.S.$ 1885. If the customer had locked in his total load of 95 MWh at the Day-ahead price of U.S.$ 20 then his total payment would have been 95 MWh * $20 = U.S.$ 1900.[5] Therefore, in this case the customer has the incentive to submit the lower demand into the day-ahead market. This result demonstrates another feature of the two-settlement system. If a participant takes a position in the opposite direction of the rest of the

---

[1]The current PJM model does not reflect the cost of marginal losses in the LMP calculation.

[2]These limitations include the generators startup time, minimum run time, energy limits, ramp limits, and other physical unit operating constraints.

[3]The PJM two settlement system consists of a day-ahead financial settlement and a real-time balancing market.

[4]This example assumes that the customer's demand is a small part of the overall market which means that the small increase in demand would not have substantially changed the clearing price.

[5]This example also assumes that the customer's demand is a small part of the overall market which means that the small increase in demand would not have substantially changed the clearing price.

market, then economic reward results. This incentive is key to driving price convergence between the markets.

The incentives illustrated in the examples encourage day-ahead market participation because if demand is too low in the day-ahead market, the customer faces higher real-time prices and if demand is too high the opposite occurs. In addition, the desire to obtain forward price certainty and to manage risk also contribute to the incentive for demand customers to submit day-ahead demand bids.

The incentives for generation are best discussed in the context of the fourth design objective. Creating the incentive for resources to follow real-time dispatch instructions is fundamental to the voluntary bid-based market design. Utilizing a voluntary market to support reliable grid operations requires strong economic signals that are consistent with real-time reliability requirements of the grid. The best way to illustrate the incentive for generation to follow real-time dispatch instructions is through the use of additional examples.

Example 3: A generator submits an incremental offer curve into the day-ahead market for 100 MW at U.S.$ 20 and for 120 MW at U.S.$ 30. The generator offer is cleared in the day-ahead market for 100 MW at U.S.$ 20. In the real-time market, the generator is requested to increase output by the system operator to the 120 MW level and the real-time LMP value is U.S.$ 31. The generator has the incentive to increase its output above the day-ahead scheduled amount because the real-time LMP is higher and the generator will receive the higher real-time price for the additional megawatt delivered in the real-time market. The generator settlement is: day-ahead $= 100$ MWh * U.S.$ 20/MWh $=$ U.S.$ 2000 and real-time $= (120-100)$ MWh * U.S.$ 31/MWh $=$ U.S.$ 620, the total generator revenue is U.S.$ 2620. This example shows that the generator has incentive to respond to real-time price increases if the real-time price exceeds the generator's incremental offer.

Example 4: A generator submits an incremental offer curve into the day-ahead market for 100 MW at U.S.$ 20 and for 120 MW at U.S.$ 30. The generator offer is cleared in the day-ahead market for 120 MW at U.S.$ 45. The next day in the real-time market, the generator is requested to reduce to 100 MW output and the real-time price falls to U.S.$ 20/MWh because of lower than expected loads. If the generator reduces output then it must purchase the quantity difference at real-time price. An examination of the generator's profit under this scenario will reveal the incentive. The Day-ahead settlement is 120 MWh * U.S.$ 45/MWh $=$ U.S.$ 5400, the generators production cost.[6] is U.S.$ 2600. If the generator does not reduce output then its profit is: U.S.$ 5400 $-$ U.S.$ 2600 $=$ U.S.$ 2800. If the generator reduces as requested, then the real-time settlement is: real-time settlement $= (100-120)$ MWh * U.S.$ 20 MWh $=$ $-$U.S.$ 400. In this case, the generator's production cost is reduced to U.S.$ 2000. The generators profit is then equal to U.S.$ 5400 $-$ U.S.$ 400 $-$ U.S.$ 2000 $=$ U.S.$ 3000. Therefore, the generator makes more profit by following the dispatch instruction to reduce output below the day-ahead scheduled level be-

[6]This example assumes that the generator submitted an offer equal to its marginal cost.
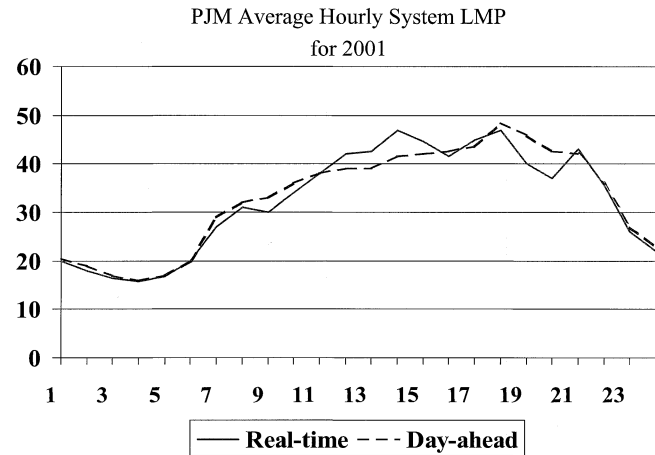

Fig. 1. Day-ahead versus real-time hourly average LMP.

cause the real-time price reduced. Essentially, the incentive is driven by the fact that the generator was paid a higher price for the day-ahead scheduled output and was able to buy it back in the real-time market at the reduced price.

This section has outlined the fundamental design concepts behind the PJM day-ahead market. The next section will examine market results.

B. PJM Day-Ahead Market Results

The PJM day-ahead energy market was implemented in June 2000. Since the implementation, the convergence between day-ahead prices and real-time prices has been narrowing [1]. In the year 2001, the average day-ahead LMP was U.S.$ 32.75/MWh and the average real-time LMP was U.S.$ 32.38. Therefore the average day-ahead LMP value was 1.1% higher than the average real-time LMP value. The relationship between the average Day-ahead LMP and the average real-time LMP varies by hour of the day. A plot of the PJM average hourly system LMP is shown in Fig. 1 [1].

The plot, Fig. 1, illustrates the close relationship that has developed between day-ahead and real-time prices. The fundamental incentives mentioned previously drive this convergence, but the convergence also occurred quickly because of the flexibility provided by the virtual bidding rules. The PJM day-ahead design makes it easy to submit virtual supply offers and demand bids which promotes liquidity and a robust market. During the year 2001, the average hourly amount of virtual supply offers cleared in the day-ahead market was 6547 MW [1]. The onpeak hour average was 8094 MW and, at times, the cleared virtual supply offers cleared were over 18 000 MW. The average hourly amount of virtual demand bids cleared in the Day-ahead market was 5393 MW [1]. The on-peak hour average was 6298 MW and, at times, the cleared virtual supply offers cleared were over 16 000 MW. These results demonstrate the high degree of liquidity in the day-ahead and forward markets.

III. REAL-TIME ENERGY MARKET

The PJM real-time energy market is based on actual real-time operating conditions. Real-time LMPs are calculated based on the actual system operating conditions as described by the PJM

state estimator using the applicable generation offer data and dispatchable external transactions. Generators that are available but not selected in the day-ahead scheduling may alter their bids for use in the real-time energy market during the generation re-bidding period from 4:00 to 6:00 P.M. (otherwise their original day-ahead market bids remain in effect for the real-time energy market). LSEs will pay real-time LMPs for any demand that exceeds their day-ahead scheduled quantities (and will receive revenue for demand deviations below their scheduled quantities). Generators are paid real-time LMPs for any generation that exceeds their day-ahead scheduled quantities (and will pay for generation deviations below their scheduled quantities). Transmission customers pay congestion charges based on real-time LMPs for bilateral transaction quantity deviations from day-ahead schedules. All spot purchases and sales in the balancing market are settled at the real-time LMPs.

A. Overview of Real-Time LMP Calculation

At 5-min intervals, locational marginal prices (LMPs) are calculated for all PJM load busses and generation busses that are modeled in the PJM state estimator. LMPs are also calculated for PJM interface busses with other control areas and for other busses outside the PJM control area as required.

In order to perform LMP calculations and to perform the associated energy settlements and billing, a complete set of input data are required. This set of input data includes

- accurate model of the actual operating conditions that exist on the PJM power grid;
- complete description of all external transactions;
- full set of offer data from generating resources;
- set of dispatchable transactions;
- list of binding transmission constraints;
- economic dispatch instructions;
- log of dispatching instructions.

The PJM state estimator provides the initial powerflow solution that is required as input to the LMP calculation.

All PJM generation that is following PJM dispatch instructions are eligible to set LMP values. In addition to generation, external transactions that are designated as dispatchable and are following PJM dispatch instructions are eligible to set LMP values. Since the LMP model includes a detailed model of adjacent external systems, loop flow effects are implicitly included in the calculations.

In the real-time energy market, PJM dispatchers meet the energy demand while respecting transmission security constraint using the least-cost security constrained economic dispatch program, the unit dispatch system (UDS). The UDS is an ex-ante dispatch that is based on the projected system conditions within the next 5 min. The LMP calculations are ex-post and are based on the actual generation response to the ex-ante dispatch that was sent 5 min before. Since the LMP calculation is based upon the actual operating conditions existing in the PJM control area as described by the PJM state estimator solution, the LMP values are calculated based on a dynamic model of the PJM power grid. The LMP calculation will take into account the current transmission and generation outages as well as transmission limitations that are identified

by the PJM dispatchers when the PJM control area is operating out of economic merit to control such limitations.

B. LMP Data Model Input

1) Offer Data From Generating Resources: Offer data from generating resources are stored in the markets database with all offer data locked as of noon the day before the actual operating day. These data along with demand bids, and external transactions are used in the PJM day-ahead market to determine the actual day-ahead scheduled quantities, day-ahead LMP values and net tie schedules. If a generator offer is accepted in the day-ahead market, then the offer carries over into the real-time market, otherwise the generator may changes its offer data for the Real-time market during the re-bid period (4:00–6:00 P.M.).

After the re-bid period, a reserve adequacy assessment is performed by PJM to determine if additional steam units, above those scheduled in the day-ahead market results, need to be scheduled in advance of the operating day.

During the operating day, the PJM dispatch will communicate desired dispatch levels by sending economic price signals (or economic dispatch rates) and/or individual unit megawatt to the generators. Generators that are following economic dispatch instructions will achieve the desired megawatt output level that is described in the dispatch signal by comparing the dispatch rate it receives from PJM to its offer curve.

2) Transaction Data: Transaction schedules can be submitted for the day-ahead energy market, day-ahead prescheduling for the real-time energy market or hour by hour into the real-time energy market during the operating day. All transaction schedules crossing the PJM control area boundary must be entered into the PJM energy management system for creation of pool-to-pool schedules and must include designation of a point of receipt and point of delivery. Points of import into the PJM control area are chosen from a set of external interface points into the PJM control area. Points of export from the PJM control area are also chosen from the set of external interface points. Points of receipt or delivery by a participant within the PJM control area will be chosen from the list of valid sources and sinks listed on the PJM OASIS. Dispatchable transactions into the PJM control area that are following economic dispatch instructions are eligible to participate in the calculation of LMPs and are modeled as generation (or load) at the designated point of receipt (or delivery).

3) Binding Transmission Constraints: When the PJM EMS system detects possible upcoming transmission limit violations, the PJM dispatch investigates solutions to the problem. If the transmission limitation can be resolved through system reconfiguration (PAR operation or switching) then the constraint is managed using the EMS system. If the transmission violation requires redispatch, the system operator transfers the transmission constraint information to the unit dispatch system for resolution. If the transmission constraint becomes a binding constraint in the security-constrained economic dispatch solution, the unit dispatch system will transfer the binding constraint information to the locational price calculation module. The LPA contingency processor translates this information and inputs the transmission constraint to the LPA. The three types of transmission constraints that can be modeled in the LPA are detailed.

1) Reactive interface limits

Voltage and stability limits that are caused by regional transfers are modeled as interface limits. Interface limits are modeled in PJM system operations, and therefore in the LPA, as a limitation to the total flow on a set of transmission lines. The PJM reactive interfaces describe key transmission boundaries on the PJM transmission system that measures the voltage performance of specific areas of the power grid.[7]

2) Thermal limits without contingency

These are limits to the amount of power that can flow on a single transmission line due to physical limitations such as conductor heating, line sag, etc. These limits are modeled in the LPA such that the flow on the line is held at or below the state estimator value.

3) Contingency limits

These limits are thermal limits that anticipate the loss of another transmission facility. A contingency limit is evaluated by calculating the resulting flow on the monitored facility for the loss of the another transmission facility (the contingency facility) or set of transmission facilities. This resulting flow is called the contingency flow. These limits are modeled in the LPA such that the contingency flow is held at or below the contingency flow that is calculated based on the state estimator solution.

C. Description of the PJM LMP Model

The PJM LMP calculation process consists of a variety of programming modules that are executed as part of the real-time sequence that executes every five minutes on the PJM energy management system (EMS). A functional diagram of the PJM LMP model is shown in Fig. 2. As indicated in Fig. 2, the main modules of the PJM LMP model are

- state estimator;
- LPA preprocessor;
- locational price algorithm (LPA);
- unit dispatch system (UDS).

Each of these modules is described in detail below. In addition to the main modules that are listed in the diagram, several other programs designed to ensure data integrity are executed as part of the PJM LMP calculation process. These programs include the LPA Input Data Consistency Check (ICC) program and the LPA output data consistency check (OCC) program.

The primary purpose of the ICC is to perform data verification on all input data to the LMP calculation process to ensure that the information is current, consistent and reasonable. The ICC program will monitor all input data files to ensure that each file's operating system timestamp and internal timestamps are current and consistent with the interval being processed. In addition, the ICC will check any transmission constraint data to verify that any contingencies entered and their corresponding controlling actions are all entered consistently, accurately, and in a timely manner. The ICC also monitors the status of the state estimator solution to ensure that the solution is a valid solved powerflow

[7]Interface limits are modeled in the LPA such that the sum of the flow on the interface will beheld at the megawatt value of the interface flow from the state estimator solution or lower.

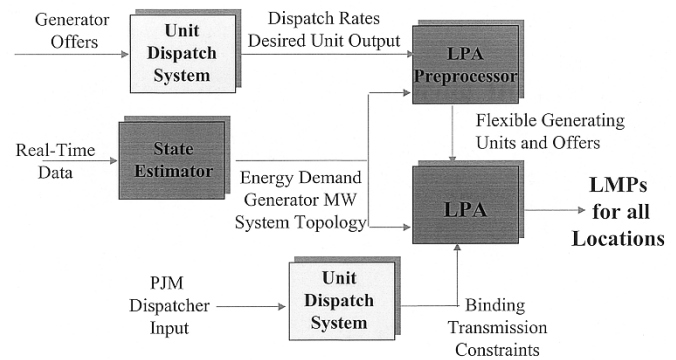### Locational Marginal Pricing Model

Fig. 2. Functional diagram of PJM LMP model.

solution. The ICC executes at the beginning of the LMP calculation sequence and if a problem is identified the program logs the error to the LPA error log and produces an appropriate alarm to the system operator.

The primary purpose of the OCC is to verify that the LMP calculation is performed accurately and completely. The OCC will check all the output data files to ensure that each program completed successfully and produced its corresponding output file along with several checks of the results for reasonability. The OCC will check the file's operating system timestamp and any internal timestamps to ensure the file's times are current and consistent with the interval being processed. In addition, the OCC will check the resultant LMP's for consistency with any constraint information. This program performs LMP validation by verifying that the LMP of each generator that is a control variable in the LPA is consistent with the price bounds that are set by the generators offer data and dispatch rate. The OCC executes at the end of the LMP calculation sequence and if a problem is identified the program logs the error to the LPA error log and produces an appropriate alarm to the system operator.

1) PJM State Estimator: The LMP calculation depends upon having a complete and coherent power flow solution as input. This input requirement can be achieved by using a state estimator. The state estimator is a standard power system operations tool whose purpose is to provide a base case power flow solution for input into other computer programs. The state estimator uses actual operating conditions that exist on the power grid (as described by metered inputs) along with the fundamental power system equations to calculate the remaining flows and conditions that are not metered. Since the state estimator solution provides a complete and coherent model of actual operating conditions based upon observable (metered) input and an underlying mathematical model, it can be used to provide the basis for the LMP calculations. The inputs to the state estimator are the available (metered) real-time measurements, the current status of equipment (lines, generators, transformers, etc.), and the bus load distribution factors.

This standard industry tool depends upon data redundancy and the underlying physical and mathematical relationships of the power system to provide a solution with less error than the

original measurements. Therefore the state estimator can correct "bad data" and calculate missing data in the model to provide a coherent representation of existing network conditions.

The PJM state estimator is run on a one minute cycle and can provide the following inputs to the LMP module

- ac powerflow solution;
- actual generator megawatt output;
- bus loads;
- tie line flows;
- megawatt losses by transmission zone;
- actual megwatt flow on any constrained transmission facility.

2) LPA Preprocessor: Since the LMP calculation is based on actual generation output rather than a theoretical optimal dispatch, it is necessary to screen generators and transactions to determine if they are eligible to participate in the LMP calculation.

The LPA preprocessor performs this screening function by analyzing the following to determine if a generator is following economic dispatch requests.

- generator state estimated megawatt hour output;
- generator offer price curves;
- economic dispatch rates;
- desired megawatt level for each generator that was specified by UDS.

The program therefore acts as a real-time performance monitoring function for generators that have designated themselves as dispatchable. Dispatchable generators whose actual megawatt output is 110% or less than the desired megawatt level are considered to be following the economic dispatch instructions and are therefore eligible to be passed through to the LMP calculation as flexible generators. The LPA preprocessor also identifies and validates any generators that are specifically requested by the UDS program to operate out of economic merit order in order to control a transmission constraint. These generators are designated as eligible to set LMP if they are on-line and following the dispatch instruction. The LPA preprocessor also screens transactions that are designated as dispatchable to determine if their offer data are consistent with current dispatch rates and if they are therefore eligible to set LMP. Generators that are not eligible to participate in LMP calculations are those that are declared must run or that are not following economic dispatch requests based on the criteria outline before.

Generators and transactions that are eligible to participate in the LMP calculation are those that are following the economic dispatch requests as described before. These eligible generators or transactions are modeled in the LMP calculation as flexible generators with offer prices that correspond to the value from their offer curve at actual megawatt output (or megawatt schedule for transactions). The LPA preprocessor calculates this real-time offer value using the following criteria

- if the generator's state estimated megawatt value is less than or equal to the desired megawatt value, then the real-time offer value is calculated by comparing the state estimated megawatt output to the offer curve.
- if the generator's state estimated megawatt value is greater than the desired megawatt value, then the real-time offer

value is calculated by comparing the desired megawatt value to the offer curve.

The eligible generators and transactions are introduced to the LPA (LPA) as flexible generators (or loads) and are modeled at actual megawatt output with a small bandwidth to allow for solution tolerance. Generators that are not eligible to participate in LMP calculations are modeled as inflexible generators with their megawatt output fixed at the actual megawatt value from the state estimator solution.

3) Locational Price Algorithm: The function of the LPA (LPA) is to determine the LMP values for each of the PJM nodes in the state estimator model and for interface busses to the PJM and PJM control areas on a 5-min basis. The LMP are defined as the cost to serve the next increment of load at each node bus location for the current system state estimated operating point and taking into account flexible generator Real-time bid prices and the buses' location with respect to transmission limitations. Given the input from the state estimator and the set of flexible generators with bid prices at their actual operating points, the calculation of LMP is relatively straight forward. The LMPrice calculation is an incremental linear optimization problem that is formulated at the current state estimated operating point. The objective is to minimize the cost function subject to the power balance constraint, generation megawatt bounds, transaction megawatt bounds and any transmission constraints that currently exist on the system.

Since the goal of the PJM LMP system is to calculate the real-time LMP values based on actual system operating conditions, the state estimated powerflow solution is used as a starting point for the incremental linear programming formulation. This powerflow solution is then linearized in order to perform the LMP calculations. As outlined in previous sections, the set of flexible generators and transactions is determined by the LPA preprocessor and is input to the LPA as a set of control variables. The set of flexible generators and sale transactions P_i and the set of flexible purchase transactions L_j are modeled in the formulation at the state estimated megawatt amount with a small bandwidth to allow for solution tolerance. The cost coefficients in the objective function for the flexible generators and transactions are set equal to the real-time bid that is calculated by the LPA preprocessor based on the state estimated megawatt level and the incremental offer (or bid) curve. These cost coefficients are assumed to have a constant slope.

The PJM incremental linear programming formulation is as follows

Minimize

$$Z = \sum C_i(\Delta P_i) - \sum C_j(\Delta L_j)$$

subject to

$$\sum \Delta P_i - \sum \Delta L_j = 0$$
$$\Delta P \min_i \leq \Delta P_i \leq \Delta P \max_i$$
$$\Delta L \min_j \leq \Delta L_j \leq \Delta L \max_j$$
$$A_{ik}\Delta P_i + D_{jk}\Delta L_j \leq 0.$$

where

ΔP_i change in power output for generator i;

$\Delta P \max_i$ upper megawatt bound for generator i;

$\Delta P \min_i$ lower megawatt bound for generator i;

C_i calculated real-time offer for generator i;

C_j calculated real-time bid for load (transaction) j;

ΔL_j change in power consumption for load j (note: in practice the only dispatchable loads that currently exit are external purchase transactions);

$\Delta L \max_j$ upper megawatt bound for load j;

$\Delta L \min_j$ lower megawatt bound for load j;

A_{ik} matrix of shift factors for generation bus i (with respect to the reference bus) on the binding transmission constraints (k);

D_{jk} matrix of shift factors for load bus j (with respect to the reference bus) on the binding transmission constraints (k).

The shift factor matrices referenced above represent a set of shift factors based on the system topology described by the binding transmission constraint. For interface limits and thermal limits w/o contingency, the transmission system configuration used to calculate the shift factors is based on the current topology as described by the state estimator. For contingency limits, the state estimator topology is modified to include the removal of the contingent transmission facilities for the shift factor calculation. Shift factors for a transmission constraint are a measure of the change in power flow on the constraint's monitored element for a unit change in megawatt injection at a bus and a corresponding unit change in megawatt withdrawal at the reference bus.

The LMP values at each bus are by-products of the linear programming formulation that is listed before. The LMP value at a particular location is essentially the sum of the marginal price of generation at the reference bus plus the marginal congestion price at the location associated with the various binding transmission constraints. The marginal prices associated with various constraints in the optimization problem are called shadow prices. The shadow price of a constraint can be explained as the incremental change in value of the objective function for a unit change in the limit (right hand side) of the constraint. Therefore, an equation for computing LMP values can be expressed in terms of these shadow prices.

The LMP equation can be written as follows:[8]

$$LMP_i = \lambda - \sum A_{ik} * SP_k$$

where

LMP_i LMPrice at bus i;

λ marginal price of generation at the reference bus;

A_{jk} shift factor for bus i on binding constraint k;

SP_k shadow price of constraint k.

It should be noted that the cost of marginal losses in the PJM LMP program is set to zero; therefore, the marginal loss term does not appear in the above equation.

4) Unit Dispatch System: The UDS is a software tool that is designed to provide the PJM dispatchers with the capability to manage changes in load, generation, interchange, and trans-

[8]The current PJM model does not reflect the cost of marginal losses in either the dispatch instructions or in the LMP values.

mission constraints simultaneously on a near real time basis, by providing a recommended dispatch solution. This solution looks at current system conditions, forecasted load, generation, and transactions to produce dispatch solutions for a user-selected look-ahead time. This allows the operator to simultaneously manage multiple transmission constraints.

The UDS is not a stand-alone system. It is an application that processes data from the markets database and other PJM systems. Other data sources include

- load forecast data from EMS (GDC servers);
- ACE, steam deviation, regulation signal from EMS (COM servers);
- constraint data-unit sensitivities from EMS (NA servers);
- state estimator output from EMS (NA servers);
- outage data from eDART;
- transaction data from EES.

The UDS executes a dispatch solution automatically every 5 min or when executed by the operator. To calculate the solution, UDS looks at online and available generation, generator bid data, forecasted load, scheduled and current interchange, area control error (ACE), and the regulation signal. The application then produces three cases with each solution. Each of these solution cases contains

- recommended set of zonal dispatch rates;
- list of exceptions to the dispatch rates for constraint control;
- individual unit dispatch rates;
- desired megawatt level for each generator.

UDS gathers transmission constraint information from the EMS alleviate overload and develops its dispatch solution based on using the most economic generators to control a given constraint. The optimization that takes place when the UDS executes will ensure that no generator is moved to alleviate one constraint, only to aggravate another constraint.

When the operator approves a recommended solution, the zonal dispatch rates and/or individual unit megawatt are bridged to the EMS, where they are automatically sent out to generators or local control centers. Units that appear on the UDS' exception list must be manually dispatched.

In addition to being used for economic dispatch, results of each recommended solution are used as input to the locational price calculation process. When the operator approves a recommended solution, dispatch rates and desired megawatt level for each generator is sent to the LPA preprocessor. UDS also produces a list on binding constraints that are sent to the LPA.

REFERENCES

[1] PJM 2001 State of the Market Rep., Norristown, PA, June 2002.

Andrew L. Ott received the B.S. degree in electrical engineering from Penn State University, University Park, and the M.S. degree in applied statistics from Villanova University, Villanova, PA.

Currently, he is responsible for the oversight of PJM market development, PJM market operations, and PJM market settlements. He joined PJM, Norristown, PA, in October 1996. He was responsible for implementation of the current PJM locational marginal pricing system, the PJM financial transmission rights auction, and the PJM day-ahead energy market. Prior to joining PJM, he was with GPU for 13 years in transmission planning and operations.