An Agent-Based Computational Laboratory for Wholesale Power Market Design

Junjie Sun and Leigh Tesfatsion, *Member, IEEE*

***Abstract*— This study reports on the model development and open-source implementation (in Java) of an agent-based computational wholesale power market organized in accordance with core FERC-recommended design features and operating over a realistically rendered transmission grid. The traders within this market model are strategic profit-seeking agents whose learning behaviors are based on data from human-subject experiments. Our key experimental focus is the complex interplay among structural conditions, market protocols, and learning behaviors in relation to short-term and longer-term market performance. Findings for a dynamic 5-node transmission grid test case are presented for concrete illustration.**

I. Introduction

In April 2003 the U.S. Federal Energy Regulatory Commission proposed the *Wholesale Power Market Platform (WPMP)* as a template for all U.S. wholesale power markets (FERC [1]). This design recommends the operation of wholesale power markets by Independent System Operators (ISOs) or Regional Transmission Organizations (RTOs) using locational marginal pricing to price energy by the location of its injection into or withdrawal from the transmission grid. Versions of this design have been implemented in New England (ISO-NE), New York (NYISO), the mid-Atlantic states (PJM), the Midwest (MISO), and the Southwest (SPP), and adopted for implementation in California (CAISO). Joskow [2, p. 6] reports that ISO/RTO operated energy regions now include over 50% of the generating capacity in the U.S.

The complexity of the WPMP market design has made it extremely difficult to undertake economic and physical reliability studies of the design using standard statistical and analytical tools. Strong opposition to the market design thus persists among some industry stakeholders due in part to a perceived lack of sufficient performance testing.

In recent years, however, powerful new agent-based computational tools have been developed to analyze this degree of complexity. The present study reports on the development and implementation of an agent-based framework for testing the dynamic efficiency and reliability of the WPMP market design. This framework – referred to as *AMES* (*A*gent-based *M*odeling of *E*lectricity *S*ystems) – models strategic traders interacting over time in a wholesale power market that is organized in accordance with core WPMP features and that operates over a realistically rendered transmission grid. To our knowledge, AMES is the first non-commercial open-source framework permitting the computational study of the WPMP design.

We are currently using the AMES framework to investigate the intermediate-term performance of wholesale power markets operating under the WPMP market design. In particular, we are exploring the extent to which this design is capable of supporting the efficient, profitable, and sustainable operation over time of existing generation and transmission facilities, despite possible attempts by some market participants to gain individual advantage through strategic pricing, capacity withholding, and induced transmission congestion.

To illustrate concretely the potential usefulness of the AMES framework for this purpose, experimental findings are reported below for a dynamic extension of a static five-node transmission grid test case used extensively for training purposes by the ISO-NE and PJM. In the static training case, the generators are assumed to report their true cost and production capacity attributes to the ISO; the possibility that generators might engage in strategic reporting behavior is not considered. In contrast, the AMES generators use reinforcement learning to decide the exact nature of the supply offers (marginal cost functions and production intervals) that they daily report to the AMES ISO for use in the WPMP day-ahead market. We show that all of the AMES generators learn over time to implicitly collude on the reporting of higher-than-true marginal costs, thus considerably raising total variable costs of operation at the ISO-determined "optimal" solutions.

Our long-run goal is to develop AMES into a framework that rings true to industry participants and policy makers and that can be used as a research and training tool. We envision academic researchers and teachers using this framework to increase their qualitative understanding of the dynamic operation of restructured wholesale power markets. Industry participants should be able to use the framework to familiarize themselves with market rules and to test business strategies. And policy makers should find the framework useful for conducting intensive experiments to explore the performance of actual or proposed market designs from a social welfare viewpoint.

II. Overview of the AMES Framework

The AMES wholesale power market framework is programmed in Java. The framework is modular, extensible, and open source in order to provide a useful foundation for further electricity research.

Last revision: December 17, 2006. To appear in the *Proceedings* of the 2007 IEEE Power Engineering Society General Meeting. This study is an abridged version of Sun and Tesfatsion [5]. This work has been supported in part by the National Science Foundation under Grant NSF-0527460.

Junjie Sun (corresponding author: jjsun@iastate.edu) and Leigh Tesfatsion (tesfatsi@iastate.edu) are with the Department of Economics, Iowa State University, Ames, IA 50011 USA.

The AMES framework currently incorporates in stylized form several core elements of the WPMP market design as implemented by the New England Independent System Operator (ISO-NE) and the Midwest Independent System Operator (MISO), respectively. By adhering closely to the architecture of these regional energy markets, we have been able to take advantage of the business practice manuals, training guides, and reports publicly released by the ISO-NE [3] and the MISO [4] for use by their market participants. These publications provide a wealth of specific implementation details missing from the more abstract WPMP template.

The core elements of the WPMP market design that have been incorporated into the AMES framework to date are as follows:

- The AMES wholesale power market operates over an AC transmission grid for DMax successive days, with each day D consisting of 24 successive hours H = $00, 01, \ldots, 23$.
- The AMES wholesale power market includes an Independent System Operator (ISO) and a collection of energy traders consisting of Load-Serving Entities (LSEs) and Generators distributed across the nodes of the transmission grid.
- The AMES ISO undertakes the daily operation of the transmission grid within a two-settlement system consisting of a Real-Time Market and a Day-Ahead Market, each separately settled by means of *locational marginal pricing*.
- During the afternoon of each day D the AMES ISO determines power commitments and *locational marginal prices (LMPs)* for the Day-Ahead Market for day D+1 based on Generator supply offers and LSE demand bids (forward financial contracting) submitted during hours $00 - 11$ of day D.
- At the end of each day D the AMES ISO produces and posts a day D+1 commitment schedule for Generators and LSEs and settles these financially binding contracts on the basis of day D+1 LMPs.
- Any differences arising during day D+1 between real-time conditions and the day-ahead financial contracts settled at the end of day D must be settled in the Real-Time Market at real-time LMPs for day D+1.
- Transmission grid congestion in the Day-Ahead Market is managed via the inclusion of congestion components in LMPs.

Five additional elements that will subsequently be incorporated into AMES to reflect more fully the dynamic operational capabilities of the WPMP market design are: (a) *market power mitigation measures*; (b) *bilateral trading*, which permits longer-term contracting; (c) a market for *financial transmission rights* to permit AMES traders to hedge against transmission congestion costs arising in the Day-Ahead Market; (d) *security constraints* incorporated into the DC OPF problems solved by the AMES ISO for the Real-Time Market and Day-Ahead Market as a hedge against system disturbances; and (e) a *(Resource Offer) Re-Bid Period* during each day D as part of a resource adequacy assessment undertaken by the AMES ISO to help ensure that forecasted loads and reserve requirements are always met. Figures 1 and 2 schematically depict the architecture and dynamic flow of this extended AMES framework.

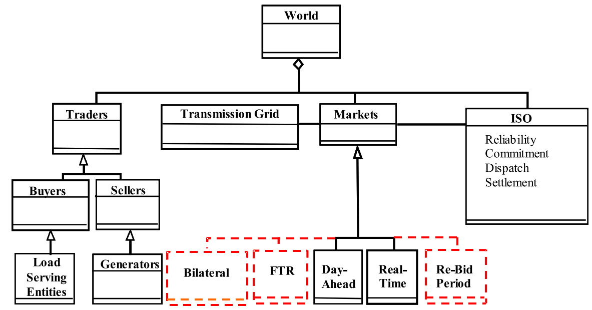

Fig. 1. AMES Architecture (Agent Hierarchy)

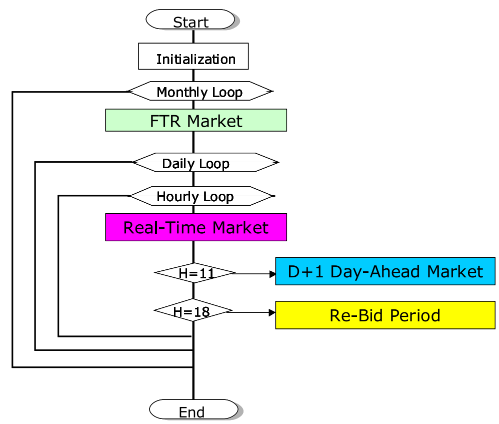

Fig. 2. AMES Dynamic Market Activities: Global View

As explained more carefully in Sun and Tesfatsion [5], the AMES ISO determines hourly power commitments/dispatch levels and LMPs for the Day-Ahead Market and Real-Time Market by solving *DC Optimal Power Flow (OPF)* problems that approximate underlying AC OPF problems. To handle these aspects, we have developed an accurate and efficient strictly convex quadratic programming (SCQP) solver module, *QuadProgJ*, wrapped in an outer DC OPF data conversion shell, *DCOPFJ* (Sun and Tesfatsion [6]). The AMES ISO solves its DC OPF problems by invoking QuadProgJ through DCOPFJ.

Trader learning is implemented in the AMES framework by a reinforcement learning module, *JReLM*, developed by Gieseler [7]. JReLM can implement a variety of different reinforcement learning methods, permitting flexible representation of trader learning within this family of methods. In later extensions of AMES, other possible trader learning methods (e.g. social mimicry and belief learning) will also be considered.

The QuadProgJ/DCOPFJ and JReLM modules for ISO grid operation and trader learning constitute the core components supporting the implementation of the AMES wholesale power market framework. This implementation is schematically depicted in Figure 3.

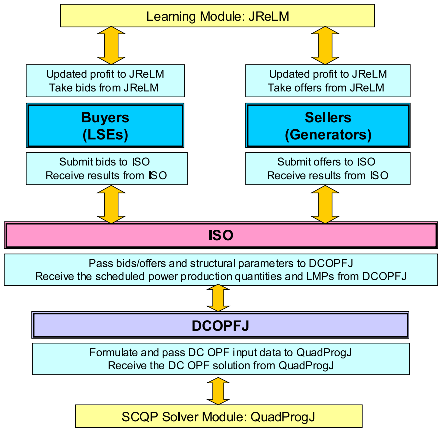

Fig. 3. Core Module Components of the AMES Framework

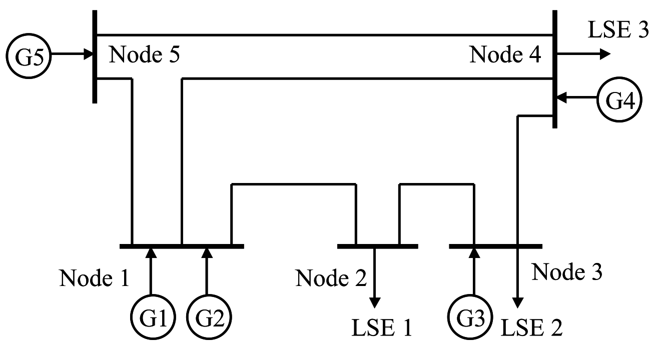

Fig. 4. A Five-Node Transmission Grid Configuration

III. Dynamic Five-Node Test Case

Consider a situation in which five Generators and three LSEs are distributed across a 5-node transmission grid as depicted in Figure 4. Originally due to John Lally [8], this five-node transmission grid configuration is now used extensively in ISO-NE/PJM training manuals to solve for DC-OPF solutions at a given point in time conditional on variously specified marginal costs and production limits for the Generators and variously specified price-insensitive loads for the LSEs. The implicit assumption in these static training exercises is that the true cost and true production limits of the Generators are known. No mention is made of the possibility that Generators in real-world ISO-managed wholesale power markets might learn to exercise market power over time through strategic reporting of their cost and production attributes.

In this section we illustrate how the AMES wholesale power market framework can be used to transform these static training exercises into a more realistic dynamic form with strategically learning Generators. Detailed grid, production, and load input data for a specific dynamic five-node test case are provided in Table I.[1]

We first ran this dynamic five-node test case under a "no learning" assumption for Generators, i.e. Generators were assumed to report to the ISO their true marginal cost functions and true production limits. Our findings for this no-learning case, reported in Sun and Tesfatsion [5], reveal the complicated effects of daily load profiles, transmission congestion, and production limits on LMP determination over time, even in the absence of strategic reporting by Generators.

We next ran this dynamic five-node test case under the assumption that the profit-seeking Generators can report strategic supply offers to the ISO. More precisely, the Generators still must report their true production limits to the ISO; but they can now learn over time what marginal cost attributes to report to the ISO in an attempt to increase their profit earnings. Using a well-known stochastic reinforcement learning algorithm explained in detail in Sun and Tesfatsion [5], each profit-seeking Generator learns over time which marginal cost function to report to the ISO based on the profits it has earned from previously reported functions.

To control for random effects, outcomes for the learning case are reported below in the form of mean and standard deviation values obtained for twenty runs using twenty different seed values. In these twenty runs, all five Generators appear to "converge" by day 422 to a sharply peaked choice probability distribution in which a probability of 0.999 is assigned to a single supply offer.[2] Consequently, all learning outcomes reported below are for day 422. Tables II and III provide detailed numerical solution values (means and standard deviations) for real power production levels and LMPs on day 422.

Figure 5 displays the (mean) solution values obtained for production for each of the 24 hours on day 422, along with the corresponding solution values obtained for day 422 in the absence of Generator learning.[3] In the no-learning case, note that the "peaker" (high cost) Generator 4 is only dispatched to produce energy at the peak load hour 17. In the learning case, however, Generator 4 is able to use strategic supply offers to ensure it is dispatched at approximately its upper production limit (200MWs) throughout each hour of the day. Also, in the no-learning case the "cheap" Generator 5 is regularly dispatched at a high production level during each hour of the day, but in the learning case it is backed way down because its strategic supply offers make it appear to be a relatively more expensive Generator. As detailed in Sun and Tesfatsion [5], this heavier reliance on costlier generation in the learning case approximately triples the total variable cost of operation.

Figure 6 graphically depicts the 24-hour (mean) LMP solution values for the learning case along with the 24-hour LMP solution values for the no-learning case. Interestingly, although the LMPs for the learning case are considerably higher than the LMPs for the no-learning case, they are also less volatile around the peak load hour 17. Consequently, the ISO is not able to use the appearance of price spikes in peak load hours to detect the considerable exercise of market power by the learning Generators. Rather, some form of direct auditing of the Generators' cost attributes would seem to be required.

[1] The transmission grid configuration, reactances, locations of the Generators and LSEs, and initial hour-0 load levels in Table I are taken from Lally [8]. The general shape of the LSE load profiles is adopted from a 3-node example presented in Shahidehpour et al. [9, p. 296-297].

[2] The *mean* convergence time across the five Generators was only 62 days.

[3] Given the stationarity of the daily load profiles and the Generators' cost functions and production limits, and the absence of system disturbances, in the no-learning case the 24-hour outcomes obtained for any one day are the same as for any other day.

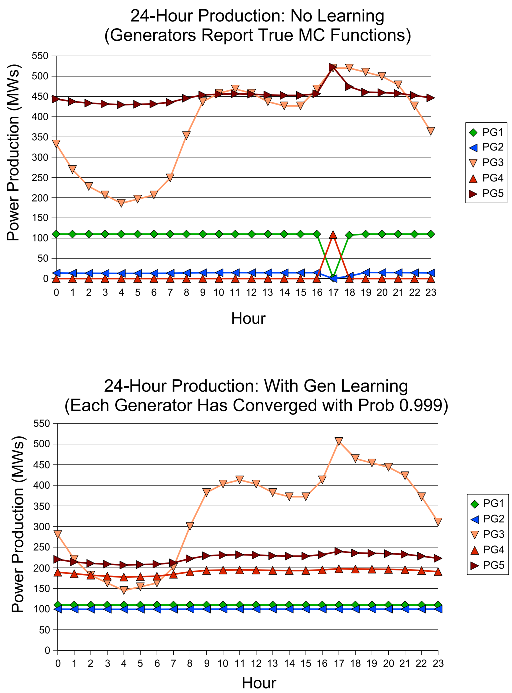

Fig. 5. Dynamic 5-Node Test Case Solution Values for 24-Hour Real Power Production Levels (Day 422) – Generator Learning Compared with No Learning

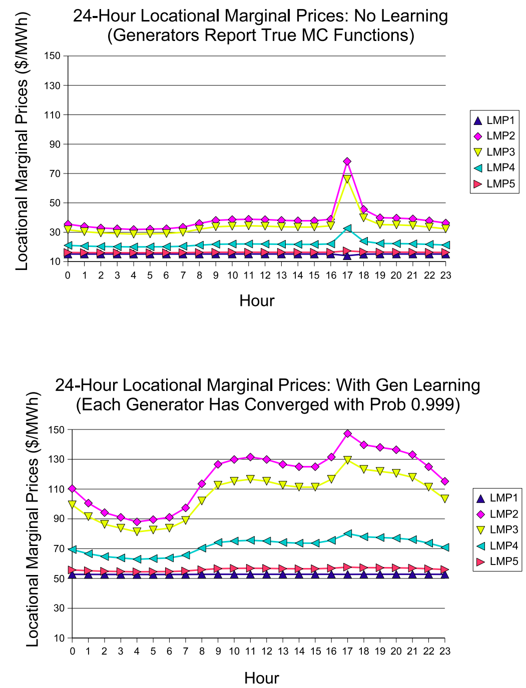

Fig. 6. Dynamic 5-Node Test Case Solution Values for 24-Hour LMPs (Day 422) – Generator Learning Compared with No Learning

Acknowledgment

We are grateful to Deddy Koesrindartoto for collaboration on earlier phases of this project, and to Daniel Kirschen, Chen-Ching Liu, James McCalley, Herman Quirmbach, Harold Salazar, and especially Anke Weidlich for helpful discussions on topics related to this paper.

References

[1] FERC, *Notice of White Paper*, U.S. Federal Energy Regulatory Commission, April 2003.

[2] P. Joskow, "Markets for power in the united states: An interim assessment," *The Energy Journal*, vol. 27, no. 1, pp. 1–36, 2006.

[3] ISO-NE. (2006) Home page. ISO New England, Inc. [Online]. Available: http://www.iso-ne.com/

[4] MISO. (2006) Home page. Midwest ISO, Inc. [Online]. Available: http://www.midwestiso.org/

[5] J. Sun and L. Tesfatsion, "Dynamic testing of wholesale power market designs: An open-source agent-based framework," Department of Economics, Iowa State University, Ames, Iowa, Working Paper No. 06025, 2006. [Online]. Available: http://www.econ.iastate.edu/tesfatsi/DynTestAMES.JSLT.pdf

[6] ——, "DC optimal power flow formulation and solution using QuadProgJ," Department of Economics, Iowa State University, Ames, Iowa, Working Paper No. 06014, 2006. [Online]. Available: http://www.econ.iastate.edu/tesfatsi/DC-OPF.JSLT.pdf

[7] C. Gieseler, "A java reinforcement learning module for the repast toolkit: Facilitating study and experimentation with reinforcement learning in social science multi-agent simulations," Department of Computer Science, Iowa State University, M.S. Thesis, 2005.

[8] J. Lally, "Financial transmission rights: Auction example," in *Financial Transmission Rights Draft 01-10-02*, m-06 ed. ISO New England, Inc., January 2002, section 6.

[9] M. Shahidehpour, H. Yamin, and Z. Li, *Market Operations in Electric Power Systems*. New York, NY: IEEE/Wiley-Interscience, John Wiley & Sons, Inc., 2002.

TABLE I

DYNAMIC 5-NODE TEST CASE – DC OPF STRUCTURAL INPUT DATA (SI)

Base Values									
S_o	V_o								
100	10								
K^a	π^b								
5	0.05								
Branch									
From	To	lineCap[c]	X^d						
1	2	250.0	0.0281						
1	4	150.0	0.0304						
1	5	400.0	0.0064						
2	3	350.0	0.0108						
3	4	240.0	0.0297						
4	5	240.0	0.0297						
Gen ID	atNode	FCost	a	b	Cap^L	Cap^U	Init$		
1	1	1600.0	14.0	0.005	0.0	110.0	$\$1M$		
2	1	1200.0	15.0	0.006	0.0	100.0	$\$1M$		
3	3	8500.0	25.0	0.010	0.0	520.0	$\$1M$		
4	4	1000.0	30.0	0.012	0.0	200.0	$\$1M$		
5	5	5400.0	10.0	0.007	0.0	600.0	$\$1M$		
LSE									
ID	atNode	L-00[e]	L-01	L-02	L-03	L-04	L-05	L-06	L-07
1	2	350.00	322.93	305.04	296.02	287.16	291.59	296.02	314.07
2	3	300.00	276.80	261.47	253.73	246.13	249.93	253.73	269.20
3	4	250.00	230.66	217.89	211.44	205.11	208.28	211.44	224.33
ID	atNode	L-08	L-09	L-10	L-11	L-12	L-13	L-14	L-15
1	2	358.86	394.80	403.82	408.25	403.82	394.80	390.37	390.37
2	3	307.60	338.40	346.13	349.93	346.13	338.40	334.60	334.60
3	4	256.33	282.00	288.44	291.61	288.44	282.00	278.83	278.83
ID	atNode	L-16	L-17	L-18	L-19	L-20	L-21	L-22	L-23
1	2	408.25	448.62	430.73	426.14	421.71	412.69	390.37	363.46
2	3	349.93	384.53	369.20	365.26	361.47	353.73	334.60	311.53
3	4	291.61	320.44	307.67	304.39	301.22	294.78	278.83	259.61

[a] Total number of nodes

[b] Soft penalty weight π for voltage angle differences

[c] Upper limit P^U_{km} (in MWs) on the magnitude of real power flow in branch km

[d] Reactance X_{km} (in ohms) for branch km

[e] L-H: Load (in MWs) for hour H, where H=00,01,...,23

TABLE II

LEARNING DYNAMIC 5-NODE TEST CASE – MEANS AND STANDARD DEVIATIONS FOR SOLUTION VALUES (SI) ON DAY 422 FOR REAL POWER PRODUCTION LEVELS (IN MWS)

Hour	$\overline{p^*_{G1}}$	p^{*SD}_{G1}	$\overline{p^*_{G2}}$	p^{*SD}_{G2}	$\overline{p^*_{G3}}$	p^{*SD}_{G3}	$\overline{p^*_{G4}}$	p^{*SD}_{G4}	$\overline{p^*_{G5}}$	p^{*SD}_{G5}
00	110.00	0.00	99.80	0.88	280.40	10.92	189.37	29.60	220.42	21.84
01	109.92	0.36	99.64	1.59	220.92	17.07	185.74	37.25	214.17	28.21
02	109.85	0.67	99.53	2.10	182.18	22.66	182.11	42.50	210.73	32.93
03	109.81	0.83	99.47	2.35	163.20	25.51	179.72	45.31	208.98	35.57
04	109.78	0.98	99.42	2.60	144.96	28.69	177.50	48.31	206.74	38.51
05	109.80	0.91	99.45	2.48	154.08	27.03	178.61	46.79	207.86	37.00
06	109.81	0.83	99.47	2.35	163.20	25.51	179.72	45.31	208.98	35.57
07	109.88	0.52	99.59	1.84	201.60	19.83	184.36	39.92	212.17	30.52
08	110.00	0.00	99.81	0.86	300.60	9.85	190.23	27.16	222.16	19.91
09	110.00	0.00	99.82	0.80	382.48	5.95	193.70	18.22	229.20	12.83
10	110.00	0.00	99.82	0.79	403.03	5.22	194.57	16.43	230.97	11.41
11	110.00	0.00	99.83	0.78	413.12	4.92	195.00	15.65	231.84	10.81
12	110.00	0.00	99.82	0.79	403.03	5.22	194.57	16.43	230.97	11.41
13	110.00	0.00	99.82	0.80	382.48	5.95	193.70	18.22	229.20	12.83
14	110.00	0.00	99.82	0.81	372.38	6.36	193.27	19.19	228.33	13.60
15	110.00	0.00	99.82	0.81	372.38	6.36	193.27	19.19	228.33	13.60
16	110.00	0.00	99.83	0.78	413.12	4.92	195.00	15.65	231.84	10.81
17	110.00	0.00	99.84	0.71	506.19	3.25	197.68	10.36	239.88	7.11
18	110.00	0.00	99.83	0.74	464.70	4.18	197.02	13.32	236.04	9.13
19	110.00	0.00	99.83	0.75	454.09	4.26	196.73	13.57	235.14	9.30
20	110.00	0.00	99.83	0.76	443.90	4.37	196.30	13.91	234.36	9.53
21	110.00	0.00	99.83	0.77	423.24	4.69	195.43	14.97	232.71	10.29
22	110.00	0.00	99.82	0.81	372.38	6.36	193.27	19.19	228.33	13.60
23	110.00	0.00	99.81	0.86	311.06	9.30	190.67	25.92	223.06	18.93

Cap^U_1	Cap^U_2	Cap^U_3	Cap^U_4	Cap^U_5
110.0	100.0	520.0	200.0	600.0

TABLE III

LEARNING DYNAMIC 5-NODE TEST CASE – MEANS AND STANDARD DEVIATIONS FOR SOLUTION VALUES (SI) ON DAY 422 FOR LMPS (NODAL BALANCE CONSTRAINT MULTIPLIERS, IN $/MWH)

Hour	$\overline{\text{LMP}_1}$	LMP^{SD}_1	$\overline{\text{LMP}_2}$	LMP^{SD}_2	$\overline{\text{LMP}_3}$	LMP^{SD}_3	$\overline{\text{LMP}_4}$	LMP^{SD}_4	$\overline{\text{LMP}_5}$	LMP^{SD}_5
00	52.74	12.33	110.30	58.16	99.39	48.02	69.40	21.56	55.70	13.06
01	52.70	12.26	100.56	49.61	91.49	41.16	66.56	19.44	55.16	12.82
02	52.68	12.23	94.18	44.34	86.32	36.92	64.69	18.12	54.81	12.67
03	52.66	12.22	91.02	41.79	83.75	34.86	63.77	17.46	54.63	12.60
04	52.63	12.23	87.96	39.38	81.27	32.90	62.86	16.84	54.45	12.54
05	52.65	12.23	89.49	40.57	82.51	33.86	63.32	17.15	54.54	12.57
06	52.66	12.22	91.02	41.79	83.75	34.86	63.77	17.46	54.63	12.60
07	52.69	12.24	97.38	46.96	88.91	39.03	65.63	18.78	54.98	12.75
08	52.75	12.37	113.52	61.04	102.01	50.33	70.35	22.28	55.87	13.15
09	52.79	12.56	126.59	73.16	112.61	60.05	74.15	25.31	56.58	13.52
10	52.80	12.62	129.87	76.28	115.27	62.55	75.11	26.09	56.75	13.61
11	52.80	12.65	131.48	77.83	116.57	63.79	75.58	26.48	56.84	13.66
12	52.80	12.62	129.87	76.28	115.27	62.55	75.11	26.09	56.75	13.61
13	52.79	12.56	126.59	73.16	112.61	60.05	74.15	25.31	56.58	13.52
14	52.78	12.53	124.98	71.64	111.30	58.83	73.68	24.93	56.49	13.47
15	52.78	12.53	124.98	71.64	111.30	58.83	73.68	24.93	56.49	13.47
16	52.80	12.65	131.48	77.83	116.57	63.79	75.58	26.48	56.84	13.66
17	52.73	12.81	147.26	92.89	129.34	75.90	80.10	30.38	57.58	14.07
18	52.80	12.81	139.68	85.72	123.22	70.13	77.95	28.50	57.26	13.93
19	52.80	12.78	138.00	84.10	121.86	68.83	77.46	28.08	57.17	13.87
20	52.80	12.75	136.38	82.54	120.55	67.58	77.00	27.68	57.09	13.82
21	52.81	12.68	133.09	79.38	117.88	65.04	76.05	26.87	56.93	13.71
22	52.78	12.53	124.98	71.64	111.30	58.83	73.68	24.93	56.49	13.47
23	52.76	12.39	115.19	62.56	103.36	51.54	70.83	22.66	55.96	13.19