Stochastic Look-Ahead Commitment: A Case Study in MISO

Preprint

Bernard Knueven,[1] M. Nazif Faqiry,[2] Manuel Garcia,[3] Yonghong Chen,[2] Roger Treinen,[4] Trevor Werho,[5] Junshan Zhang,[5] Vijay Vittal,[5] Long Zhao,[2] Anupam Thatte,[2] Shengfei Yin[2]

1 National Renewable Energy Laboratory
2 Midcontinent Independent System Operator
3 Sandia National Laboratories
4 Nexant, Inc
5 Arizona State University

Presented at the 2023 IEEE Power & Energy Society General Meeting
Orlando, Florida
July 16-20, 2023

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Conference Paper
NREL/CP-2C00-84539
September 2023

Stochastic Look-Ahead Commitment: A Case Study in MISO

Preprint

Bernard Knueven,[1] M. Nazif Faqiry,[2] Manuel Garcia,[3] Yonghong Chen,[2] Roger Treinen,[4] Trevor Werho,[5] Junshan Zhang,[5] Vijay Vittal,[5] Long Zhao,[2] Anupam Thatte,[2] Shengfei Yin[2]

1 National Renewable Energy Laboratory
2 Midcontinent Independent System Operator
3 Sandia National Laboratories
4 Nexant, Inc
5 Arizona State University

Suggested Citation
Knueven, Bernard, M. Nazif Faqiry, Manuel Garcia, Yonghong Chen, Roger Treinen, Trevor Werho, Junshan Zhang, Vijay Vittal, Long Zhao, Anupam Thatte, Shengfei Yin. 2023. *Stochastic Look-Ahead Commitment: A Case Study in MISO: Preprint*. Golden, CO: National Renewable Energy Laboratory. NREL/CP-2C00-84539. https://www.nrel.gov/docs/fy23osti/84539.pdf.

© 2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Conference Paper
NREL/CP-2C00-84539
September 2023

National Renewable Energy Laboratory
15013 Denver West Parkway
Golden, CO 80401
303-275-3000 • www.nrel.gov

NOTICE

This work was authored by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by the U.S. Department of Energy Advanced Research Projects Agency – Energy. The views expressed herein do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free via www.OSTI.gov.

Cover Photos by Dennis Schroeder: (clockwise, left to right) NREL 51934, NREL 45897, NREL 42160, NREL 45891, NREL 48097, NREL 46526.

NREL prints on paper that contains recycled content.

Stochastic Look-Ahead Commitment: A Case Study in MISO

Bernard Knueven, M. Nazif Faqiry, Manuel Garcia, Yonghong Chen, Roger Treinen, Trevor Werho, Junshan Zhang, Vijay Vittal, Long Zhao, Anupam Thatte, Shengfei Yin

***Abstract*—This paper introduces the Stochastic Look Ahead Commitment (SLAC) software prototyped and tested for the Midcontinent Independent System Operator (MISO) look ahead commitment process. SLAC can incorporate hundreds of wind, load, and net scheduled interchange (NSI) uncertainty scenarios. It uses a progressive hedging method to solve a novel two-stage stochastic unit commitment. The first stage commitment decisions, made only for those generators whose decision to commit or not in each time period cannot be deferred, can cover the uncertainties within the next three hours. The second stage includes both the dispatch for each of the scenarios and the commitment decisions that can be deferred. Study results on 15 MISO production days show that SLAC may bring economic and reliability benefits under uncertainty.**

***Index Terms*--Stochastic Optimization, Unit Commitment, Uncertainty Management**

I. Introduction

With the growing levels and corresponding uncertainty and variability of stochastic resources, along with associated changing operating conditions and more frequent extreme weather events, power system operators are faced with significant challenges in operating the system securely. Existing tools are based on deterministic optimization models that consider a small range of scenarios. This paper introduces a prototype advisory tool — Stochastic Look-Ahead Commitment (SLAC) — that may potentially be used by system operators to enhance system security and improve energy market surplus under growing uncertainty and variability.

Like many ISO/RTOs, the amount of renewable energy resources is growing significantly within the Midcontinent Independent System Operator (MISO) footprint, increasing the amount of uncertainty that grid operators must manage. Today, ISO/RTOs use deterministic clearing engines, offline studies, and statistical analysis of historical data for a subset of inputs. For example, headroom margin is applied in the forward Reliability Assessment Commitment (RAC) and Look Ahead Commitment (LAC) processes to allow additional capacity to handle input data uncertainties.

SLAC leverages statistical information from an ensemble of potential operational scenarios and their respective likelihood. For a given time period, SLAC will calculate an optimal solution (e.g., commitment and dispatches) over the study period that maximizes the *expected* market surplus over all the operational scenarios it considers. References [1, 2] introduced how MISO operations currenlty manage uncertainties through a multi-stage commitment process, reserve products, "headroom" and multiple scenarios. LAC runs every 15 minutes with three hours look ahead, and is the last stage of the commitment process. A robust look ahead commitment was prototyped in 2013 [2] and showed potential operational benefit. However, it also indicated computational challenges.

Unlike most stochastic unit commitment problems studied in the literature [3, 4, 5, 6], a typical SLAC problem has relatively few committable resources available (most units are committed by the day-ahead market), studies a shorter time-horizon with higher fidelity (three hours with 15-mintue windows), and is re-solved frequently (every 15 minutes). Because of these factors, the formulation we choose defers those commitment decisions it can (due to notification time), to a later instance of the problem, when more of the uncertainty for a given time period has been resolved. Therefore, the "stages" in the SLAC formulation are defined as "those commitment decisions which *cannot* be deferred" and "those commitment decision which *can* be deferred," regardless of the time period in which these decisions happen.

II. Formulation

First, consider a generic stochastic optimization problem [7, 8]:

$$\text{minimize } \ (c \cdot x) + \sum_{s \in S} p_s (\, d_s \cdot y_s)$$
$$\text{subject to: } \ (x, y_s) \in Q_s, \ \ \forall s \in S$$

which minimizes some (linear) objective in expectation over every scenario s, given here-and-now decisions x, which must be *identical* for every scenario s, and wait-and-see decisions y_s, which can be *different* in each scenario s. Each scenario s may take on a unique probability $p_s > 0$, with the condition $\sum_{s \in S} p_s = 1$. In the context of SLAC, the here-and-now

B. Knueven is with the National Renewable Energy Laboratory, Golden, CO, 80401 USA. M. N. Faqiry, Y. Chen, L. Zhao, A. Thatte, and S. Yin are with the Midcontinent Independent System Operator, Inc. (MISO), Carmel, IN, 46032 USA. M. Garcia is with Sandia National Laboratories, Albuquerque, NM, 87185 USA. R. Treinen is with Nexant, Inc., Chandler, AZ, 85226 USA. J. Zhang and V. Vittal are with Arizona State University, Tempe, AZ, 85287 USA.

This work was supported by the United States Department of Energy under the Grant DE-AR0000696. This work was authored in part by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

Disclaimer: The views expressed in this paper are solely those of the authors and do not necessarily represent those of MISO.

decisions are generator commitments in time periods that must be decided on in this SLAC run — that is, commitments in time periods which will become fixed decisions in the SLAC run 15 minutes in the future due to the notification time requirements of the individual generators.

A. Notification Times and Non-Anticipatory Constraints

Each individual generator has a specific notification time, which represents the amount of time that generator requires to be notified before changing its commitment status. SLAC fixes each generator's commitment status to its current status for the duration of the generator's notification time. In the time period immediately following the exhaustion of a generator's notification time, SLAC enforces the commitment decision in all scenarios to be identical, becoming part of the first-stage decisions x. After each generator's "first committable" time period, the commitment decisions are deemed second-stage variables (e.g., part of y_s for each scenario). These later commitment decisions can be deferred until at least the next SLAC run, 15 minutes later, when more uncertainty has been resolved. Therefore, the first-stage decisions made by the SLAC become "final" because of the generator's notification time they are fixed in the very next SLAC run.

To summarize, as opposed to most stochastic unit commitment models, the proposed SLAC has here-and-now decisions x which are determined purely by *individual* generators' notification time, as opposed to fixed time periods. SLAC defers to the second stage decisions y_s all commitments which could be finalized fifteen-minutes into the future or later. In this fashion, SLAC represents the flexibility inherit in the system by allowing different commitments in different scenarios; only those commitments which need to be finalized in the next fifteen minutes are enforced to be non-anticipative.

B. Scenario Generation

The scenario generation procedure is described in reference [8]. For MISO, this procedure develops probabilistic forecasts of the power output of wind farms in its footprint, the load from 37 local balancing authorities, and the total net-scheduled-interchange with neighbors. The probability distributions capture the spatial and temporal correlations between the uncertainties from these three sources. The probabilistic forecasts are sampled 200 times and are evenly weighted. A backwards reduction method based on the Kantorovich distance between pairwise scenarios is applied iteratively until the we arrive at 40 probabilistically weighted scenarios. The scenarios are updated during every 15-minute interval that the SLAC is solved.

C. LAC Operational Constraints

The operational constraints defined by the set Q_s are modeled to match MISO LAC requirements and rules pertaining to resources, transmission network, reserve zones, and overall MISO system [9]. In practice, Q_s is mixed-integer linear representable. On a resource level, Q_s includes on/off and start-up/shut-down constraints, generator limits and ramp rate constraints, minimum up/down time constraints, reserve (e.g., reg, spin, supplemental) provision constraints, and some commitment-fixing constraints to honor prior and subsequent commitment plans. On the network level, Q_s includes activated transmission and post-contingency reserve deployment constraints imposed on a predetermined set of transmission constraints fed from production [10]. The remaining constraints in Q_s relate to system-level power balance and reserve requirements.

The system-level reserve requirements in MISO LAC are satisfied through reserve constraints for regulation, contingency, and ramp capability. These constraints are implemented to hold capacity and accommodate real-time load fluctuations, contingencies, and ramp needs due to net load volatility stemming from variability and uncertainty in demand, wind generation, net scheduled interchange (NSI), and other uncertainties in the input data. The SLAC tool implements these reserve requirements in each LAC-equivalent model that is built for each uncertainty scenario. While there is the potential for a stochastic look-ahead commitment to implicitly as opposed to explicitly model some of these reserve requirements [11], all the reserve products considered in SLAC are to maintain appropriate operational flexibility at a sub-fifteen-minute timescale.

D. Model Validation

The LAC optimization problem with the operational constraints Q_s were modeled in *Pyomo* [12, 13] in a customized version of *Egret* [14, 15]. The Egret model was validated and benchmarked against MISO LAC by solving and comparing an extensive number of small and large test case problems including a set of 1,436 full-size LAC problems from 15 MISO operation days each containing ~1,200 generators and numerous transmission constraints. The benchmarking criteria were to obtain the same or close enough results (e.g., objective value, commitment, energy/reserve schedule, and transmission flow/violation) when comparing MISO LAC and Egret solutions. MISO's LAC is one of the largest existing practical look-ahead unit commitment problems with sophisticated market rules and implementation logic. Extensive benchmarking against MISO's LAC and practical implementations of the SLAC on large complex MISO test cases is an important novelty in this study that distinguishes it from existing research applied on academic test cases.

III. SOLUTION METHODOLOGY

A. Progressive Hedging Implementation

We utilize the progressive hedging (PH) [16], implementation available as part of the *mpi-sppy* package [7], which operates on existing deterministic-equivalent Pyomo models to create equivalent stochastic optimization formulations and supporting algorithms. We note that PH has been applied with success on small- to medium-scale systems for day-ahead or longer unit commitment considering load and/or renewables uncertainty [17, 18, 19, 20, 21], though to our knowledge this is its first published uses for an industrial-scale stochastic unit commitment problem. The mpi-sppy package allows the modeler to significantly customize the PH algorithm for both performance and practical usability; SLAC makes use of subproblem grouping or bundling, generator-

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

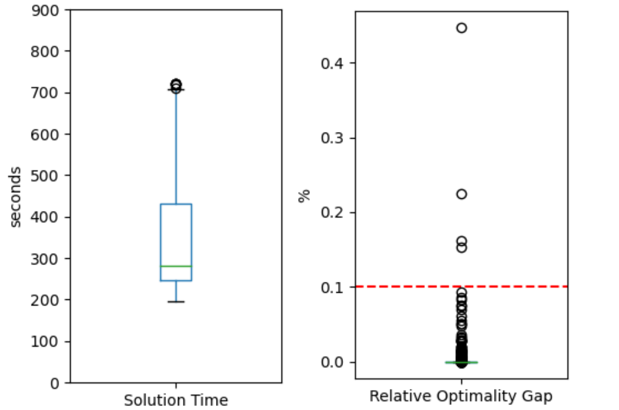

Fig. 1. Box-and-whiskers plots demonstrating SLAC solution time in seconds (left) and SLAC solution quality as measured by the computed relative optimality gap (right).

specific values of the PH step-size parameter ρ, and an incumbent solution recovery heuristic, which we now describe.

1) Solution recovery and iteration limit

We set a PH iteration limit of 10, and if PH has not converged in 10 iterations, we construct a non-anticipative solution as follows. Recall each generator has at most a single first-stage commitment decision, associated with the time period LAC is first able to change its status. Let x be a vector of these first stage 0/1 commitment decisions associated with each generator, i.e., $x_g = 1$ means generator g is committed and $x_g = 0$ means generator g is *not* committed. Consider $\bar{x}^k$, the average commitment across all scenarios, which at termination is not necessarily 0/1 valued. We then construct a first stage solution $\hat{x}$ as follows. If $\bar{x}_g^k > 0$ for generator g, then we set $\hat{x}_g \coloneqq 1$ (i.e., we commit the generator), conversely, if $\bar{x}_g^k = 0$ for generator g, then we set $\hat{x}_g \coloneqq 0$ (i.e., we do not commit the generator). On days with reliability issues (e.g., reserve shortfall), this simple heuristic outperformed other approaches – the worst optimality gap over our entire test set was 0.5% – by committing any resource still needed in any scenario after 10 PH iterations.

2) SLAC Solver Solution Times

We tracked the quantitative performance, both in terms of wall-clock time and solution quality, of each individual SLAC problem solved over 15 days' worth of rolling horizon simulations for a total of 1,436 individual SLAC optimization problems. This gives a broad perspective of the computational performance of the SLAC solver, spanning different days, seasons, system configurations, and system stressors.

All computational evaluations were completed on a virtual machine provided by MISO, with 32 virtual CPUs and 256GB RAM. All SLAC solver runs used 20 concurrent threads and the CPLEX 20.1. Reported times are wall-clock times.

As can be seen in Fig. 1, all 1,436 SLAC instances are solved well-within the 15-minute (900 second) time limit established by MISO, and the majority are solved within five minutes (300 seconds). The solution time results in Fig. 1 include the time to read data from the disk, set up all Pyomo models, execute the PH algorithm, compute the objective value from the solution recovery heuristic, and write the full scenario

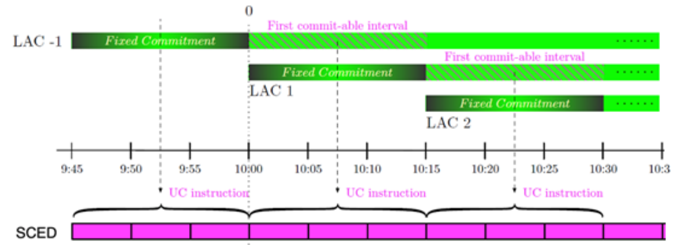

Fig. 2. Interaction between LAC/SLAC and SCED. Three consecutive LAC/SLAC iterations are shown. LAC/SLAC decide commitment statuses used in the SCED and the SCED determines the initial conditions used in the next iteration of the LAC/SLAC.

solutions to the disk. It should be noted that these reported times are not overly optimized: better optimization of the Pyomo model build-time could halve the set-up time. Further, computing and writing the full SLAC solution (including recourse) is not strictly necessary for executing the here-and-now decision. Based on results from open-source stochastic unit commitment problems run at NREL, we would also expect further returns to parallelism with more compute resources [7].

3) SLAC Solver Solution Quality

We measure the solution quality using the typical "relative optimality gap" measure. That is, the relative optimality gap gap, for a given solution UB and given lower bound LB is $gap = (UB - LB)/UB$. We set a target solution quality, or relative optimality gap, of 0.1% for the SLAC, consistent with general practice. In Fig. 1, we detail in a box-and-whiskers plot the SLAC solution quality obtained over the 1,436 instances examined. As seen, all but four out of 1,436 SLAC problems (99.7%) are solved to a 0.1% optimality gap, and most (over 96%) meet a 0.01% optimality gap requirement. Finally, we note that this is a conservative estimate of the bound, as we obtained the lower bound LB at no additional computational cost during PH execution. With more computational resources, a potentially better lower bound could be computed from each PH iteration [22], or other approaches for generating strong lower bounds could be applied [23]. Regardless, in every case the computed relative optimality gap was less than 0.5%, which is generally be considered an acceptable solution.

IV. Rolling Horizon Simulations

A. Overview of the Rolling Horizon

In this section we use the rolling horizon simulations to compare the performance of three different models of the LAC. The first is the deterministic LAC using the MISO forecast used in practice (termed MISO LAC). The second is the deterministic LAC using the point forecast developed in [15], which provides a forecast using a stochastic model that considers past MISO forecast performance as well as the relationship between past measured values and future power outputs (termed ASU LAC) [8]. The third is the SLAC with 40 scenarios. We use a few simplifications in the rolling horizon simulations in this section. First, the SCED intervals are assumed to be solved every fifteen minutes as opposed to the five-minute frequency used in practice. The interaction between the LAC and the single period real-time SCED is illustrated in Fig. 2. Second, we assume that the system

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

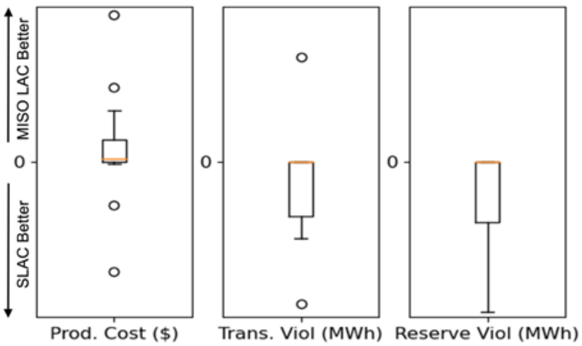

Fig. 3. Box and whiskers plot illustrating the relative total production costs, transmission constraint violations, and reserve constraint violations for each of the 15 days studied. Each plot represents the SLAC value relative to the MISO LAC value. A positive value indicates the value is larger for the SLAC and a negative value indicates the value is higher for the MISO LAC.

operator follows the commitment decisions provided by the LAC/SLAC. Indeed, in practice the LAC and SLAC are only advisory tools, and the system operator may choose not to implement their suggested commitment decisions. Finally, we neglect the forecast update between the LAC/SLAC and the SCED, assuming that the same forecast information is available when solving the LAC/SLAC and the SCED.

Overall, we see significant improvements utilizing SLAC over the two deterministic LAC models during days where the system is particularly stressed, otherwise the SLAC performs similarly to the two deterministic LAC models. This is perhaps in contrast with what might be expected with a stochastic model — in our experiments SLAC does not significantly increase production cost unless a reliability benefit is also observed. When SLAC provides improvements over LAC we typically observe these improvements in one of two ways. The first and most common observed improvement are increases reliability by reducing constraint violations: SLAC reduces transmission constraint violations and reserve constraint violations as compared to the deterministic LAC models but increases production costs by a small amount. The second type of improvement we observe are decreases in production costs as compared to the deterministic LAC models: the SLAC solution significantly decreases production costs as compared to the deterministic LAC solutions while maintaining similar constraint violation levels.

B. Rolling Horizon Simulation Results

Simulations were conducted for 15 days throughout 2018 and 2019. These days were chosen to represent stressed and conservative operation days from different seasons throughout the year with possible constraint violations and higher production costs. Fig. 3 provides a box and whiskers plot that illustrates the relative total production cost, transmission violations and reserve violations for each of the 15 days studied. The plots represent the SLAC value relative to the MISO LAC value and a positive value indicates that the SLAC value is larger. Due to MISO data confidentially, only the relative values (not absolute values) are shown.

As is common, Fig. 3 illustrates that the SLAC production costs are typically slightly larger than the production costs resulting from the MISO LAC; however, this increase in production cost is typically small and is on the order of 0.01% of the total production cost for the system. Furthermore, outliers in the relative production cost tend to be symmetric around the mean. In other words, large savings in production cost are about equally as likely when using the SLAC or MISO LAC. On average, we conclude that the production costs do not significantly favor either the MISO LAC or the SLAC.

The transmission and reserve constraint violation plots in Fig. 3 illustrate that the benefits of SLAC are primarily realized by avoiding these constraint violations. None of the 15 days exhibited more reserve constraint violations when using the SLAC and only one of the 15 days exhibited more transmission constraint violations when using the SLAC (and on this day SLAC was able to avoid reserve violations instead). Furthermore, we occasionally see significant reductions in transmission and reserve violations when using the SLAC – these reductions in transmission and reserve violations are typically associated with higher production costs.

Fig. 4 illustrates the production cost and constraint violations for five characteristic days that were chosen to illustrate five distinctly different outcomes of using the SLAC as compared to deterministic LAC models. As compared to LAC, we observe days where SLAC performs similarly (Day 1), reduces transmission violations (Day 2), reduces reserve violations (Day 3), shifts higher-cost reserve violations to lower-cost transmission violations (Day 4), and reduces production costs (Day 5). Throughout the 15 days studied, we never see a day in which the SLAC performs worse than the LAC; however, we do see some days where the SLAC and LAC perform very similarly.

V. CONCLUSIONS

This paper introduces the SLAC software prototyped and tested for the MISO LAC process. The tool utilizes a novel two-stage stochastic unit commitment formulation, where the non-anticipative decisions are driven by generator notification times. The computational performance with the customized PH method demonstrates the feasibility of real-world application. Rolling horizon simulations illustrate the performance benefit of using SLAC versus alternative deterministic LAC formulations. SLAC is shown to improve system reliability by reducing transmission and reserve constraint violations with little impact on production costs.

The main limitation of our results follows from the fact that we are using historical data from 2018 and 2019. First, it is difficult to disentangle manual operator actions, such as emergency deployment and operator over-rides, from the historical data. To the extent that manual operator actions resolve unexpected events, we underestimate the value of SLAC, because the SLAC is intended to resolve these unexpected events itself. Second, these results are limited to the uncertainties that existed in the MISO system in the years 2018 and 2019. Indeed, we expect SLAC would likely become more valuable in the future if the net-demand uncertainty increases from an increasing penetration of uncertain renewable generation. Third, the forecasts used in this study are only intended to capture uncertainty in net-demand and

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

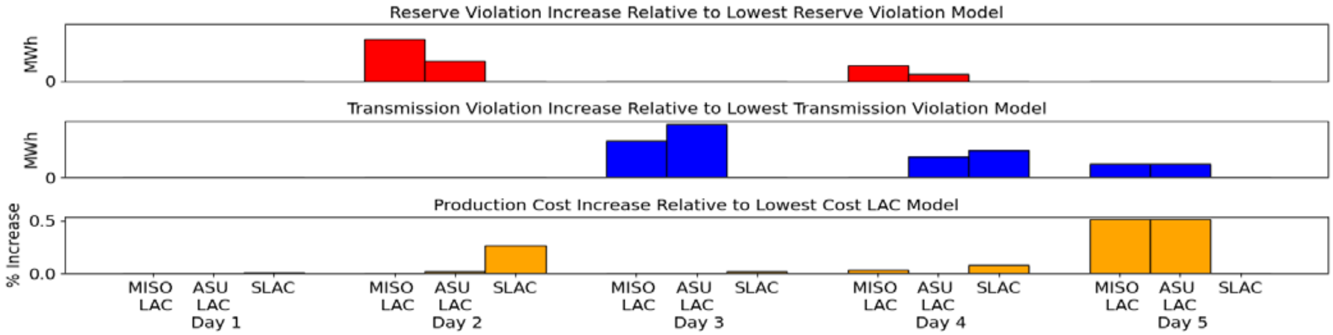

Fig. 4. Three bar plots comparing the performance of our three models for the first group of five days. For a given day the bottom bar plot shows the difference in production cost between each of the three models and the lowest cost model. The middle and top plots are similar. The middle plot represents the difference in transmission violation and the top plot represents the difference in reserve violation.

NSI. As a result, these forecasts are not intended to capture the many other uncertainties that an ISO/RTO experiences in practice. Future studies could extend this by additionally forecasting uncertainties not captured in our work.

Acknowledgment

The authors thank Jessica Harrison, John Harmon, Congcong Wang and Jason Howard from MISO for their support and valuable discussions.

References

[1] Y. Chen, V. Ganugula, J. Williams, J. Wan and Y. Xiao, "Resource transition model under MISO MIP based Look Ahead Commitment," in *2012 IEEE Power and Energy Society General Meeting*, 2012.

[2] Y. Chen, Q. Wang, X. Wang and Y. Guan, "Applying robust optimization to MISO Look-Ahead commitment," in *2014 IEEE PES General Meeting | Conference & Exposition*, 2014.

[3] Q. P. Zheng, J. Wang and A. L. Liu, "Stochastic Optimization for Unit Commitment—A Review," *IEEE Transactions on Power Systems,* vol. 30, no. 4, pp. 1913-1924, 2015.

[4] M. A. H. Håberg, "Fundamentals and recent developments in stochastic unit commitment," *International Journal of Electrical Power & Energy Systems,* vol. 109, pp. 38-48, 2019.

[5] K. Kim and V. M. Zavala, "Large-Scale Stochastic Mixed-Integer Programming Algorithms for Power Generation Scheduling," , pp. 493-512, 2016.

[6] M. Tahanan, W. v. Ackooij, A. Frangioni and F. Lacalandra, "Large-scale Unit Commitment under uncertainty," *A Quarterly Journal of Operations Research,* vol. 13, no. 2, pp. 115-171, 2015.

[7] B. Knueven, D. Mildebrath, C. Muir, J. D. Siirola, J.-P. Watson and D. L. Woodruff, "A Parallel Hub-and-Spoke System for Large-Scale Scenario-Based Optimization Under Uncertainty," Optimization Online, 2020. [Online]. Available: http://www.optimization-online.org/DB_FILE/2020/11/8088.pdf.

[8] T. Werho, J. Zhang, V. Vittal, Y. Chen, A. Thatte and L. Zhao, "Scenario Generation of Wind Farm Power for Real-Time System Operation," arXiv, 16 June 2021. [Online]. Available: https://arxiv.org/abs/2106.09105.

[9] Midcontinent Independent System Operator, Inc., "BPM 002 - Energy and Operating Reserve Markets," [Online]. Available: https://www.misoenergy.org/legal/business-practice-manuals/.

[10] Y. Chen, P. Gribik and J. Gardner, "Incorporating post zonal reserve deployment transmission constraints into energy and ancillary service co-optimization," *IEEE Trans. on Power Syst.,* vol. 29, no. 2, pp. 537-549, March 2014.

[11] B. Wang and B. F. Hobbs, "Real-Time Markets for Flexiramp: A Stochastic Unit Commitment-Based Analysis," *IEEE Transactions on Power Systems,* vol. 31, no. 2, pp. 846-860, 2016.

[12] M. L. Bynum, G. A. Hackebeil, W. E. Hart, C. D. Laird, B. L. Nicholson, J. D. Siirola, J.-P. Watson and D. L. Woodruff, Pyomo--optimization modeling in python, Third ed., vol. 67, Springer Science & Business Media, 2021.

[13] W. E. Hard, J.-P. Watson and D. L. Woodruff, "Pyomo: modeling and solving mathematical programs in Python," *Mathematical Programming Computation,* vol. 3, no. 3, pp. 219-260, 2011.

[14] M. Bynum, A. Castillo, C. Laird, B. Knueven and J.-P. Watson, "EGRET: Electrical Grid Research and Engineering Tools," 2021. [Online]. Available: https://github.com/grid-parity-exchange/Egret.

[15] B. Knueven, J. Ostrowski and J.-P. Watson, "On Mixed-Integer Programming Formulations for the Unit Commitment Problem," *Informs Journal on Computing,* vol. 32, pp. 857-876, 2020.

[16] R. T. Rockafellar and R. J. B. Wets, "Scenarios and policy aggregation in optimization under uncertainty," *Mathematics of Operations Research,* vol. 16, no. 1, pp. 119-147, 1991.

[17] S. Takriti, J. Birge and E. Long, "A stochastic model for the unit commitment problem," *IEEE Transactions on Power Systems,* vol. 11, no. 3, pp. 1497-1508, 1996.

[18] S. M. Ryan, R. J.-B. Wets, D. L. Woodruff, C. Silva-Monroy and J.-P. Watson, "Toward scalable, parallel progressive hedging for stochastic unit commitment," in *2013 IEEE Power & Energy Society General Meeting*, 2013.

[19] C. Ordoudis, P. Pinson, M. Zugno and J. M. Morales, "Stochastic unit commitment via Progressive Hedging — extensive analysis of solution methods," in *2015 IEEE Eindhoven PowerTech*, 2015.

[20] B. Rachunok, A. Staid, J.-P. Watson, D. L. Woodruff and D. Yang, "Stochastic Unit Commitment Performance Considering Monte Carlo Wind Power Scenarios," in *2018 IEEE International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)*, 2018.

[21] K. Cheung, D. Gade, C. Silva-Monroy, S. M. Ryan, J.-P. Watson, R. J.-B. Wets and D. L. Woodruff, "Toward scalable stochastic unit commitment. Part 2: Solver Configuration and Performance Assessment," *Energy Systems,* vol. 6, no. 3, pp. 417-438, 2015.

[22] D. Gade, G. Hackebeil, S. M. Ryan, J.-P. Watson, R. J.-B. Wets and D. L. Woodruff, "Obtaining lower bounds from the progressive hedging algorithm for stochastic mixed-integer programs," *Mathematical Programming,* vol. 157, no. 1, pp. 47-67, 2016.

[23] N. Boland, J. Christiansen, B. C. Dandurand, A. C. Eberhard, J. T. Linderoth, J. R. Luedtke and F. Oliveira, "Combining Progressive Hedging with a Frank--Wolfe Method to Compute Lagrangian Dual Bounds in Stochastic Mixed-Integer Programming," *Siam Journal on Optimization,* vol. 28, no. 2, pp. 1312-1336, 2018.

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.