Near-Optimal Solutions for Day-Ahead Unit Commitment

Dominic Yang*, Bernard Knueven†, Jean-Paul Watson‡, James Ostrowski§
* Mathematics and Computer Science Division
Argonne National Laboratory, Lemont, IL 60439 USA
† Computational Science Center
National Renewable Energy Laboratory, Golden, CO 80401 USA
‡ Center for Applied Scientific Computing and Global Security Directorate
Lawrence Livermore National Laboratory, Livermore, CA 94550 USA
§ Industrial and Systems Engineering
University of Tennessee Knoxville, Knoxville, TN 37996 USA

***Abstract*—Given the difficulty and the time pressure of solving unit commitment problems, near-optimal solutions (those with 0.1 or 0.001% optimality gaps) are often used in practice. The choice in which of the near-optimal solutions is used, however, is random. We investigate the impact of solution choice on the revenues obtained by generator owners across a variety of pricing schemes and problem instances.**

***Index Terms*—Mixed-integer linear programming (MILP), unit commitment (UC), market-clearing, symmetry.**

Nomenclature

Indices and Sets

$g \in \mathcal{G}$	Thermal generators.
$t \in \mathcal{T}$	Hourly time steps: $1, \ldots, \mathbf{T}$.

Parameters

$c^g(p)$	Cost function for generator g to generate p MW of power ($/MWh).
$\mathbf{C}_U^g$	Cost coefficient for generator g to start up ($/MWh).
$\mathbf{C}_D^g$	Cost coefficient for generator g to shut down ($/MWh).
$\mathbf{D}_t$	Load (demand) at time t (MW).
$\mathbf{DT}^g$	Minimum down time for generator g (h).
$\overline{\mathbf{P}}^g$	Maximum power output for generator g (MW).
$\underline{\mathbf{P}}^g$	Minimum power output for generator g (MW).
$\mathbf{RD}^g$	Ramp-down rate for generator g (MW/h).
$\mathbf{RU}^g$	Ramp-up rate for generator g (MW/h).
$\mathbf{SD}^g$	Shutdown ramp rate for generator g (MW/h).
$\mathbf{SU}^g$	Start-up ramp rate for generator g (MW/h).
$\mathbf{UT}^g$	Minimum up time for generator g (h).

Variables

p_t^g (P_t^g)	Power output for (representative) generator g at time t (MW).
u_t^g (U_t^g)	Commitment status of (rep) g at time t, $\in \{0,1\}$.
v_t^g (V_t^g)	Startup status of (rep) g at time t, $\in \{0,1\}$.
w_t^g (W_t^g)	Shutdown status of (rep) g at time t, $\in \{0,1\}$.

Funding provided by U.S. Department of Energy Office of Electricity's Advanced Grid Modeling Program. 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. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. This work was performed in part under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. The authors would like to thank Gurobi for providing the academic solver license for executing all computational experiments. This work was supported by the U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research, under Contract DE-AC02-06CH11357.

I. Introduction

Day-ahead energy market clearing in the United States is typically performed based on some unit commitment solution (UC) obtained from a commercial mixed-integer program (MIP) solver. Because it is desirable to have a UC solution within 10 to 15 minutes and proving the optimality of a solution is often computationally difficult, if the solver can certify optimality within a certain percentage – usually 0.1%, the system operators will terminate the day-ahead MIP. Unfortunately, this can lead to schedules that are unfair: a more expensive generator may be scheduled over a cheaper generator. This can clearly be the case when a generator is small enough that the cost of its dispatch is within the optimality gap, but it can also occur when the difference in cost between two generators is within the optimality gap.

The paper undertakes a study of UC solutions "in the gap," that is, additional solutions that are within a specified optimality gap. We propose a method for enumerating a diverse set of UC solutions – exploiting the symmetry of inherit in the data – to consider solutions that are within the typical optimality gap used by industry – 0.1% or 0.01%. Hence all of the solutions considered could be classified as "high quality" but can result in significantly different distribution of revenue to participating generators.

II. Methodology

To study UC solutions "in the gap," we need to have the capability to enumerate a large class of UC solutions, preferably ones which are sufficiently distinct from each other. Modern commercial Mixed-Integer Program (MIP) solvers allow the user to generate potentially all near-optimal solutions for a given problem by adjusting parameters (in Gurobi this is accomplished by adjusting the `PoolGap` and `PoolSolutions` parameters). However, as noted in [1], this may be computationally prohibitive as the number of near-optimal solutions can be large, and even if that is not the case, proving there are no solutions within a certain `PoolGap` may still take a significant amount of computational effort. The methods described next reformulate the problem such that a single solution in the MIP solver can potentially represent many, many solutions to the original UC problem – in one case in our tests, we represented over 10^{15} distinct UC solutions a a single solution to our symmetry-aware UC reformulation.

The main source of trouble in enumerating solutions is (near) symmetry induced by generators having the same (or similar) technical parameters. Two generators that are identical can have their schedules, or even parts of their schedules, swapped with no change in objective value. Similarly, two generators with similar physical characteristics and objective function can be similarly swapped with little change in objective value. Even if two similar generators are not co-located at the same bus, their schedules may still be interchangeable if the part of the transmission network linking them is sufficiently non-congested. In the context of the wholesale electricity market, such solutions distribute revenue differently to energy producers. Because of this, participants might win or lose out on significant profit depending on the near-optimal solution the optimization engine happened upon first. For more on the structural symmetry in UC, the reader is referred to [2]–[4].

A. Symmetry-Aware Unit Commitment Formulation

We discribe next a simple version of the well-known "3-bin" implementation from the literature that models the problem using three binary variables encoding the UC, start-up, and shutdown status for each generator at each hour in the event horizon. This simplified version of the model does not take into account reserves, costs for cold or hot starts, or the network. For these details the reader is referred to the "tight" formulation from [5].

$$\min \sum_{g\in\mathcal{G}}\sum_{t\in\mathcal{T}} c^g(p_t^g)u_t^g + \mathbf{C}_U^g v_t^g + \mathbf{C}_D^g w_t^g \tag{1a}$$

subject to:

$$\sum_{g\in\mathcal{G}} p_t^g = \mathbf{D}_t \qquad \forall t\in\mathcal{T} \tag{1b}$$

$$\underline{\mathbf{P}}^g u_t^g \le p_t^g \le \overline{\mathbf{P}}^g u_t^g \qquad \forall g\in\mathcal{G}, \forall t\in\mathcal{T} \tag{1c}$$

$$p_t^g - p_{t-1}^g \le \mathbf{RU}^g u_{t-1}^g + \mathbf{SU}^g v_t^g \qquad \forall g\in\mathcal{G}, \forall t\in\mathcal{T} \tag{1d}$$

$$p_{t-1}^g - p_t^g \le \mathbf{RD}^g u_t^g + \mathbf{SD}^g w_t^g \qquad \forall g\in\mathcal{G}, \forall t\in\mathcal{T} \tag{1e}$$

$$u_t^g - u_{t-1}^g = v_t^g - w_t^g \qquad \forall g\in\mathcal{G}, \forall t\in\mathcal{T} \tag{1f}$$

$$\sum_{i=t-\mathbf{UT}^g+1}^{t} v_i^g \le u_t^g \qquad \forall g\in\mathcal{G}, \forall t\in\mathcal{T} \tag{1g}$$

$$\sum_{i=t-\mathbf{DT}^g+1}^{t} w_i^g \le 1 - u_t^g \qquad \forall g\in\mathcal{G}, \forall t\in\mathcal{T} \tag{1h}$$

$$p_t^g \in \mathbb{R}_+,\ u_t^g, v_t^g, w_t^g \in \{0,1\} \qquad \forall g\in\mathcal{G}, \forall t\in\mathcal{T}. \tag{1i}$$

For simplicity, we also do not explicitly state the form of the cost function for a given generator to generate p MW of power $c_g(p)$. In practice, we model this as a piecewise linear function.

One method to address the problems induced by symmetry in MIPs is that of orbital shrinking, introduced in [6] and analyzed in the context of UC in [3]. For generators that have identical (or at least very similar) parameters, we can consider equivalence classes of generators $\mathcal{E} = \{[g] | g \in \mathcal{G}\}$, where we define an equivalence relation by stating that $g \sim g'$ if and only if g and g' have identical parameters. If there are any identical generators, we can reduce the size of this problem by introducing general integer variables for each equivalence class indicating how many generators are on and replacing the binary variables for each generator; [3] describes conditions under which this can be done exactly. Using capital letters to denote integer variable version of each binary variable and any equivalence class by $\mathcal{E}_g$ where g is the representative for the class and $|\mathcal{E}_g|$ is the size of the class, we can write our new problem as the following MIP:

$$\min \sum_{\mathcal{E}_g\in\mathcal{E}}\sum_{t\in\mathcal{T}} c^g(P_t^g)U_t^g + \mathbf{C}_U^g V_t^g + \mathbf{C}_D^g W_t^g \tag{2a}$$

subject to:

$$\sum_{\mathcal{E}_g\in\mathcal{E}} P_t^g = \mathbf{D}_t \qquad \forall t\in\mathcal{T} \tag{2b}$$

$$\underline{\mathbf{P}}^g U_t^g \le P_t^g \le \overline{\mathbf{P}}^g U_t^g \qquad \forall t\in\mathcal{T}, \forall\mathcal{E}_g\in\mathcal{E} \tag{2c}$$

$$P_t^g - P_{t-1}^g \le \mathbf{RU}^g U_{t-1}^g + \mathbf{SU}^g V_t^g \qquad \forall t\in\mathcal{T}, \forall\mathcal{E}_g\in\mathcal{E} \tag{2d}$$

$$P_{t-1}^g - P_t^g \le \mathbf{RD}^g U_t^g + \mathbf{SD}^g W_t^g \qquad \forall t\in\mathcal{T}, \forall\mathcal{E}_g\in\mathcal{E} \tag{2e}$$

$$U_t^g - U_{t-1}^g = V_t^g - W_t^g \qquad \forall t\in\mathcal{T}, \forall\mathcal{E}_g\in\mathcal{E} \tag{2f}$$

$$\sum_{i=t-\mathbf{UT}^g+1}^{t} V_i^g \le U_t^g \qquad \forall t\in\mathcal{T}, \forall\mathcal{E}_g\in\mathcal{E} \tag{2g}$$

$$\sum_{i=t-\mathbf{DT}^g+1}^{t} W_i^g \le |\mathcal{E}_g| - U_t^g \qquad \forall t\in\mathcal{T}, \forall\mathcal{E}_g\in\mathcal{E} \tag{2h}$$

$$P_t^g \in \mathbb{R}_+,\ U_t^g, V_t^g, W_t^g \in [0, |\mathcal{E}_g|]_{\mathbb{Z}} \qquad \forall t\in\mathcal{T}, \forall\mathcal{E}_g\in\mathcal{E}. \tag{2i}$$

It is not difficult to come up with problems where solutions to the aggregate model cannot be decomposed into feasible solutions to the disaggregated model. However, having a large set of generators gives the system some additional flexibility, and using a tight formulation for the ramping constraints (see [5]) can often ameliorate this problem. In practice, for the systems we consider, almost all of the generators have redundant ramping constraints when considered with an hourly time horizon.

B. Counting Disaggregated Solutions

The motivation for using formulation (2) is that it encodes some or most of the symmetry in the problem – therefore single solution to problem (2) can represent many solutions to the original UC problem (1). In some cases this reformulation means that the solver can tractably enumerate all near-optimal solutions to (2). For a given solution $\{U_t^{g*}, P_t^{g*}\}$ to problem (2), we can separate the solution by equivalence classes $\mathcal{E}$. For each $\mathcal{E}_g \in \mathcal{E}$, the number of solutions represented by (U^{g*}, P^{g*}) can be calculated by enumerating all solutions to the following auxiliary system of equations:

$$\sum_{g \in \mathcal{E}_g} p_t^g = P_t^{g*} \qquad \forall t \in \mathcal{T} \tag{3a}$$

$$\sum_{g \in \mathcal{E}_g} u_t^g = U_t^{g*} \qquad \forall t \in \mathcal{T} \tag{3b}$$

$$\text{constraints (1b)} - \text{(1i)}. \tag{3c}$$

When aggregation/disaggregation can be done exactly, the number of solutions to problem (1) represented by a single solution $\{U_t^{g*}, P_t^{g*}\}$ to problem (2) can be calculated by enumerating the solutions to problem (3) for each $\mathcal{E}_g \in \mathcal{E}$. For $\mathcal{E}_g \in \mathcal{E}$, call the number of solutions to problem (3) $N(\mathcal{E}_g)$. Then the number of solutions to problem (1) represented by a single aggregate solution $\{U_t^{g*}, P_t^{g*}\}$ is $\prod_{\mathcal{E}_g \in \mathcal{E}} N(\mathcal{E}_g)$.

We were able to calculate that some instances from the CAISO generator set (described in Section III-A) have over 10^{15} solutions within 0.001% of optimal. As this is a lower bound on the number of near-optimal solutions, analyzing all near-optimal solutions (for any reasonable tolerance) would surely be intractable.

C. Computing Representative Solutions

To account for this issue, we construct a set of near-optimal solutions that are representative of the collection of near-optimal solutions in that they reflect the diversity of the solution set but not its size. We first construct a set of aggregate solutions that should implicitly describe the symmetries of the system, and then we exploit these symmetries to build a diverse set of near-optimal solutions.

First, consider the "no-good" cut for a given solution $\{u_t^{g*}\}$:

$$\sum_{u_t^g : u_t^{g*}=1} (1 - u_t^g) + \sum_{u_t^g : u_t^{g*}=0} u_t^g \geq d. \tag{4}$$

This has the effect of rendering the given UC solution infeasible and requiring any new solution to be at least hamming distance d away from the given solution. If $d = 1$, then the solution needs to differ only at one point. This formulation works only for binary variables, but it can be extended to work with general integer variables under a binarization scheme.

Enumerating all the near-optimal solutions to the aggregated problem (2) can still be computationally intractable. Additionally, for each aggregate solution, one can have a huge number of disaggregated solutions, and explicitly computing all of them (even using a simple metric like hamming distance) can easily result in an intractable problem. To sidestep both issues, we exploited some advanced features typically present in commercial MIP solvers, namely, the ability to add lazy constraints on the fly. We also put an absolute cap on the number of solutions to limit the overall scale of the problem.

We generated these diverse near-optimal solutions in two passes. In the first pass, we use the aggregated problem (2), to construct valid lower and upper bounds for U_t^g for each $\mathcal{E}_g$ across a large set of aggregated solutions, either to a specific optimality tolerance ϵ or to a maximum number of aggregated solutions N. To enumerate all near-optimal solutions, we repeatedly solve (2) with a no-good cut (as a lazy constraint) of the form (4) with $d = 1$ for all solutions found. We then use the minimal and maximal values of U_t^g across all found near-optimal solutions as lower and upper bounds. If our search terminates before hitting the solution limit N, then the upper and lower bounds are valid for all solutions within ϵ of the optimal solution. If not, then these are heuristic bounds on U_t^g across a large number of near-optimal solutions.

In the second pass, we construct a representative set of disaggregated solutions. To do so we add the constraints (5) to the disaggregated problem (1), where LB_t^{g*} and UB_t^{g*} are the computed lower and upper bounds, respectively.

$$LB_t^{g*} \leq \sum_{g' \in \mathcal{E}_g} u_t^{g'} \leq UB_t^{g*} \qquad \forall \mathcal{E}_g \in \mathcal{E} \tag{5}$$

Then we repeatedly solve the amended problem, adding a no-good cut with hamming distance d for the found solution after each solve. We repeat until the problem is infeasible, ultimately finding a diverse set of disaggregated solutions, each pair hamming distance d apart.

D. Computing Prices

We compute the generator profits and revenues for each UC solution under three different pricing schemes from [7], all of which use some linear programming analogue of (1) and compute prices as the dual value of the power balance constraints (1b).

The first pricing scheme considered is the Approximate Convex Hull Price (aCHP), which uses the linear programming relaxation of (1). While not commonly used in practice, convex hull pricing in electricity markets has been a subject of research for several years [8]–[11]. In this context, the word "approximate" refers the fact that we have only approximate convex hulls for each generator in the pricing problem, as opposed to a full convex hull description. Large (though not exponential) convex hulls for a large classes of generators are known [12], [13], but we will not consider those here.

The remaining two pricing schemes are more common in practice: Enhanced Locational Marginal Price (ELMP) and Locational Marginal Price (LMP). The ELMP prices-in the start-up and no-load costs of the scheduled units by fixing u_t^g to 0 if generator g is not scheduled in time t, and otherwise relaxes it to be in $[0, 1]$. Hence, the price reflects only the fixed and marginal costs for the generators that are scheduled, a potential drawback for aCHP. The LMP pricing scheme fixes

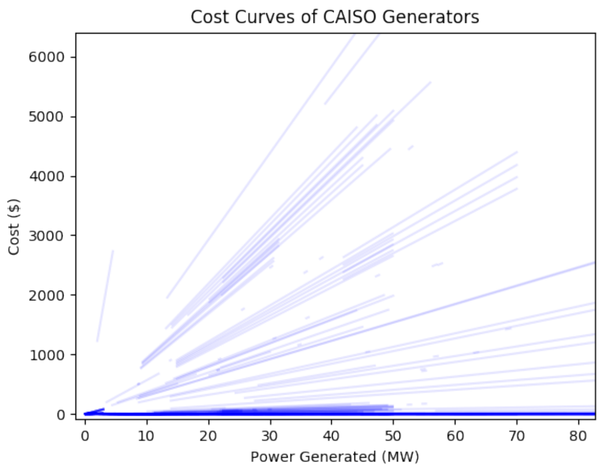

Fig. 1. Cost as a function of power generated plotted for a select number of CAISO generators. Each cost curve is plotted only on the range of possible power generation for a given generator

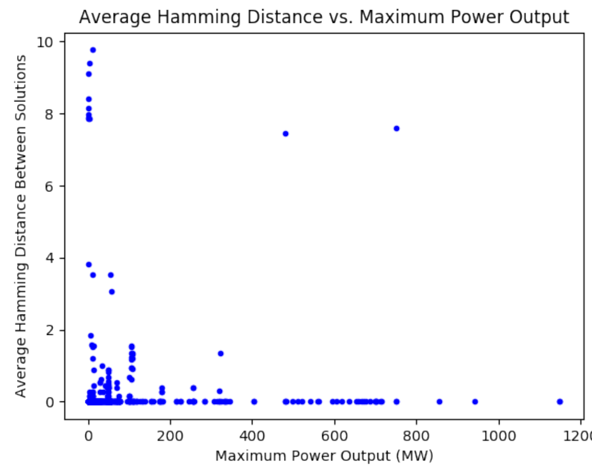

Fig. 2. Average hamming distance between each pair of solutions as a function of generator maximum power output. This data refers to the September 1 2014 instance. These are all CAISO generators.

each generator to its commitment from the UC solution, and in this way the LMP only reflects the marginal cost for the generators that are scheduled. For a more through overview of these pricing schemes, the reader is again referred to [7].

III. Computational Experiments

We applied our method to two fleets of generators. The first set of generators to be considered is the California ISO (CAISO) set of generators. CAISO manages the majority of California's power operation. The second under consideration is the Reliability Test System Grid Modernization Lab Consortium (RTS-GMLC) collection modeled using a hypothetical forecasted demand and weather profile in 2020.

A. Caiso Instances

The CAISO fleet of generators consists of 610 generators, the vast majority of which are gas and biomass engines. This class of generators is characterized by having a large number of generators with very low power generation capacity with half of the generators having a maximum power output of less than 30 MW.

This set of generators has a small degree of exact symmetry with subsets of identical generators ranging from sizes 2 to 8. When aggregating generators that are exactly the same, it is possible to reduce the complexity to approximately 3/4 the size of the original problem with 466 equivalence classes of generators that are exactly the same.

Given the high amounts of symmetry apparent in the generator specifications as well as the preponderance of small generators, we should expect a high amount of near-optimal solutions that can be produced by interchanging generators in equivalence classes. This is what occurs and we have seen in the tests that at least 10000 solutions were produced within a MIP Gap of 0.01% and in the majority of cases, over 100000 could be produced.

We explored instances using demand profiles selected from 4 dates beginning in September 2014 up until June 2015, and with reserve levels at 0, 1, 3, and 5 percent. We used 48 hours of demand data starting midnight on the date of a given instance, and produced 48-hour commitment schedules. In each of these instances, we asserted that solutions be within 0.01% of the optimum and that they exhibited pairwise hamming distances of at minimum 100 in their UC vectors. The high level of symmetry was again made apparent: in all instances, we found at minimum 12 solutions mutually of 100 hamming distance separated across all 20 instances.

Generators that actually assumed different UC schedules tended to be small generators with low maximum power outputs. For evidence of this fact, observe the clear negative trend where average hamming distance decreases as power output in Figure 2. It appears that variation in solutions occurs in small generators perhaps because their operation (or lack thereof) has a minimal effect on optimal cost. That being said, there typically are two or more large generators that effectively flip schedules. Across all near-optimal solutions, we observed rigidity in the scheduling in that over all instances tested, roughly three-fourths of the generators (on average 464 out of 610) did not exhibit any differences in their schedules across the computed solutions. Of those that did exhibit change, the typical differences were more commonly changes corresponding to turning generator off or on slightly earlier or later.

We compiled revenues and profits at a per-generator level for each solution where energy was priced using the aCHP, LMP, and ELMP schemes. A plot of the average standard deviation of the revenues under each of these schemes is shown in Figure 3.

As can be observed in the plots of average standard deviation, the ELMP scheme exhibits the greatest variation, often by one order of magnitude, over the other pricing schemes,

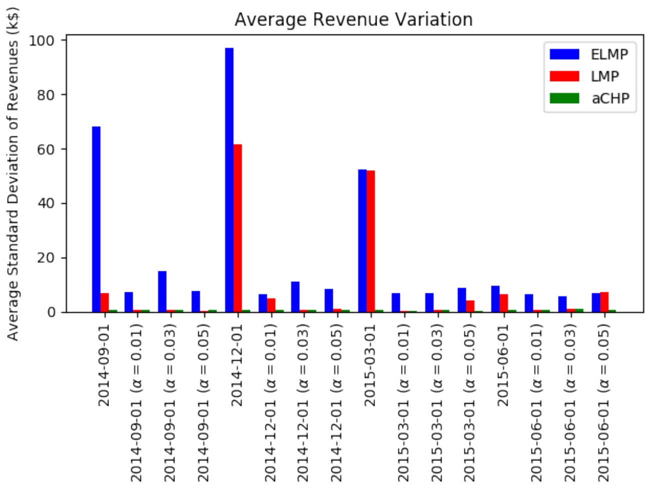

Fig. 3. Averages of the standard deviations of the revenues across near-optimal solutions in the CAISO power system according to various pricing schemes.

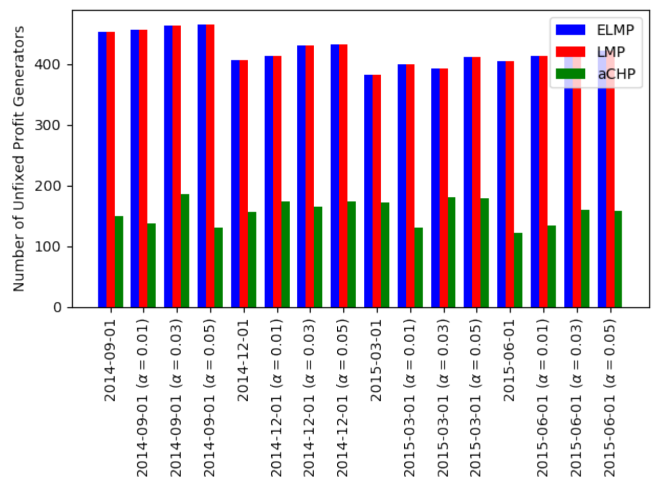

Fig. 4. Number of generators in each CAISO instance that show a difference in revenue generated across solutions.

followed by the LMP scheme and then by aCHP scheme which exhibits the least variation. For both LMP and ELMP, these price spikes are primarily caused by the scheduling of a small, under the margin, generator with a high marginal cost of production, which for the UC solution was still near-optimal (0.01%) due to the generator's size relative to the system.

We see that in all cases the ELMP pricing scheme exhibits the greatest variation and the LMP scheme generally exhibits the second-most variation although there are a few cases where the aCHP scheme shows higher variation. The higher levels of volatility in the ELMP and LMP schemes can be in part explained by the fact that even if a generator has exactly the same schedule across all solutions, it can produce different revenues depending on the schedules of the other generators. The aCHP pricing scheme lacks this freedom since the scheme depends only on the integer program constraints and not on the exact solution chosen. The difference can be observed in the bar plot in Figure 4 which depicts the number of "unfixed revenue generators", that is, the number of generators that exhibit any difference in revenue across near-optimal solutions. The amount of unfixed aCHP revenue generators is exactly the number of generators that show a difference in the UC schedule whereas the number of unfixed LMP or ELMP revenue generators is almost triple that number.

B. RTS-GMLC Instances

The RTS-GMLC collection of generators is a much smaller set than the CAISO fleet, consisting of only 73 thermal generators, approximately half being gas generators and the other half being split fairly evenly between steam and coal generators with the addition of 1 nuclear generator. The fleet contains a fairly diverse set of generators, with a group of about 50 generators with capacities below 100 MW and smaller groups of generators with larger capacities at about 150 MW and 350 MW. The collection of thermal generators is supplemented by a plethora of solar panels and hydro generators which often amount to 40% of the demand on a given day. The RTS-GMLC case also distributes its generators on a network that introduces additional constraints onto the model. We will consider both the case where the generators are all located on a single bus without the effect of network constraints (which we refer to as "CopperSheet") and the case with the network included.

Given the much smaller set of generators, there are fewer opportunities for exact and partial symmetry, yet there are still quite a few clusters of sizes ranging from 2 to 6 of generators with the exact same specifications. Joining these generators into equivalence classes reduces the complexity from 73 generators to 39 classes with an average 1.87 generators per cluster. When introducing the network, only generators on the same bus can be joined together and so the number of generators per cluster decreases further.

We considered instances a month apart in the year 2020 where we used the forecasted demand and renewable energy to produce a well-posed problem that we solved to a MIP gap of 0.1%. With the small set of generators, it is computationally tractable to enumerate all solutions for a giving instance. We enumerated all such solutions and counted the maximal solution set size, where all solutions are pairwise a given hamming distance apart. These counts for the CopperSheet case are given in Table I. Note that in certain instances, there is exactly 1 solution within the desired MIP gap. The case with the network included produced even fewer solutions because there are more constraints on the model.

Even with such a sparse set of solutions, there still exist large deviations in the possible revenue and profit generated for each given generator. With the high penetration of renewable energy in the power grid, the possibility of demarginalization exists wherein renewable energy is not fully dispatched, causing the price (in all three schemes) to drop to zero (since renewable energy has zero marginal cost in our model). This means that many generators are operating at a complete loss but cannot be turned off because the shutdown

TABLE I
NUMBER OF SOLUTIONS FOR THE RTS-GMLC COPPERSHEET INSTANCE THAT ARE PAIRWISE A GIVEN HAMMING DISTANCE APART. NOTE THAT FOR HAMMING DISTANCE 1, THIS IS JUST THE NUMBER OF SOLUTIONS TO THE UNIT COMMITMENT PROBLEM.

Instance	1	10	50	100
2020-01-31	214.0	2.0	2.0	1.0
2020-03-01	4.0	2.0	2.0	1.0
2020-03-31	126.0	1.0	1.0	1.0
2020-04-30	2.0	2.0	2.0	1.0
2020-05-30	11.0	1.0	1.0	1.0
2020-06-29	2.0	1.0	1.0	1.0
2020-07-29	4.0	1.0	1.0	1.0
2020-08-28	1.0	1.0	1.0	1.0
2020-09-27	1012.0	43.0	3.0	1.0
2020-10-27	1.0	1.0	1.0	1.0
2020-11-26	2.0	2.0	1.0	1.0
2020-12-26	3.0	3.0	1.0	1.0

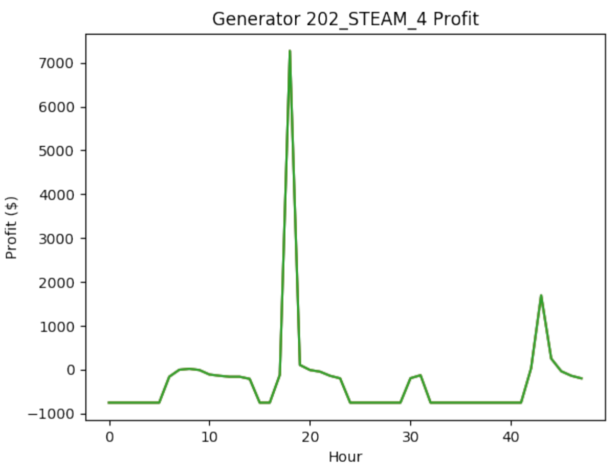

Fig. 5. Hourly profits for 3 solutions for a specific generator in the 2020-01-30 instance in the RTS GMLC CopperSheet power system. The generator 202_STEAM_4 is scheduled to run for all 48 hours but for certain stretches of time is operating at a complete loss due to demarginalization. This was computed using the aCHP pricing scheme.

cost is comparable to running the generator at a loss for 48 hours. An example of this can be seen in Figure 5.

We computed revenues and profits on a per-generator basis for each near-optimal solution computed in each instance. Differences appear primarily due to deciding whether one generator should be shut down for an entire day or kept on, and the bulk of the costs comes from the hour of shut down or start up. This behavior can be seen in the UC schedule and hourly profits in Figures 6 and 7. In the profit plot, we observe that on startup it incurs a big cost due to starting up, and only at peak power usage times do we see large profits.

In assessing the difference between pricing schemes, we note that there are not large disparities between the different pricing schemes. As for the CAISO instances, we plot the

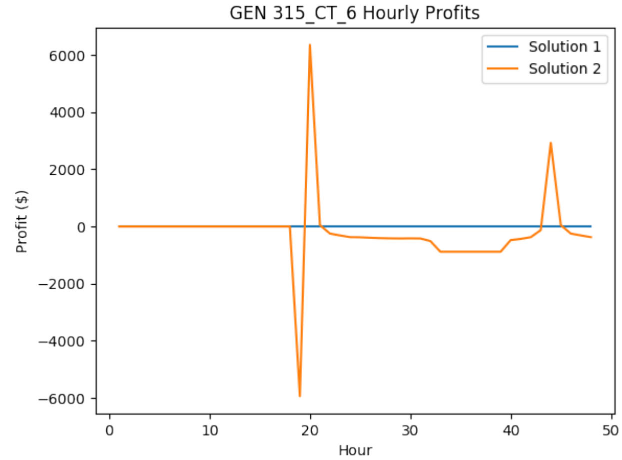

Fig. 6. Hourly profits in the aCHP pricing scheme on a hourly basis for generator 315_CT_6 in the RTS GMLC CopperSheet system for two solutions to the 2019-09-27 instance.

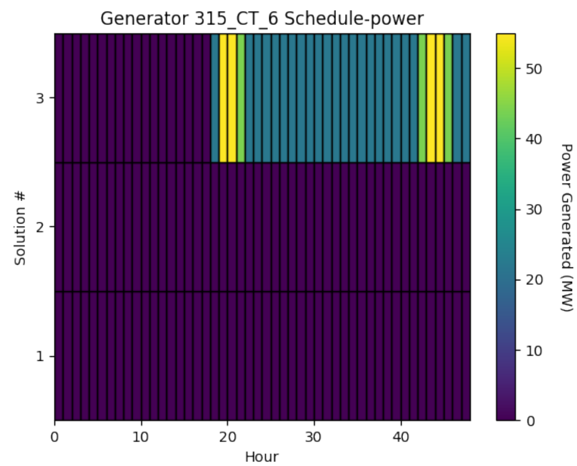

Fig. 7. Power generation on a hourly basis for generator 315_CT_6 in the RTS GMLC CopperSheet system for three solutions to the 2019-09-27 instance.

average over all generators of the standard deviations of the computed near-optimal solutions. This is displayed in Figure 8 for the CopperSheet case and Figure9 for the network. We observe that in both network and non-network cases, the variation among the pricing schemes is very close. Broadly, we observe that the ELMP appears more variable than the aCHP, which appears more variable than the LMP.

We suspect that the small differences between the pricing schemes are due to the very low number of solutions in consideration and furthermore, the low number of generators. In the non-network case, we observed on average that approximately 10 generators exhibited any difference in power produced, and in the network case, only an average of 6 generators exhibited any difference. Since the variation in revenue occurred primarily due to flipping the schedule of two generators, the difference in revenue might be explained

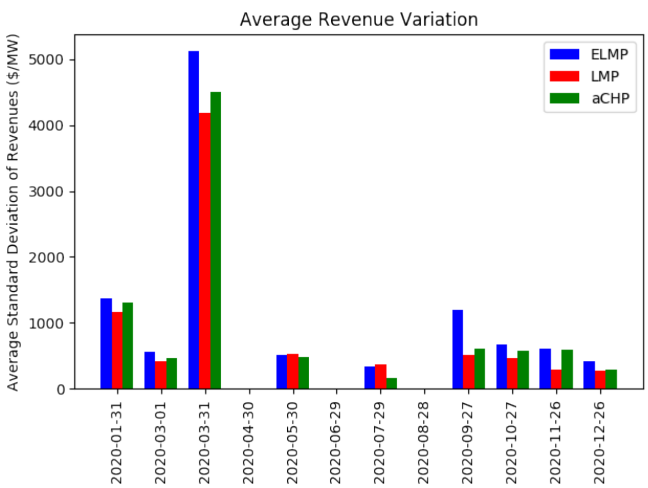

Fig. 8. Averages of the standard deviations of the revenues across near-optimal solutions in the RTS GMLC CopperSheet power system according to various pricing schemes.

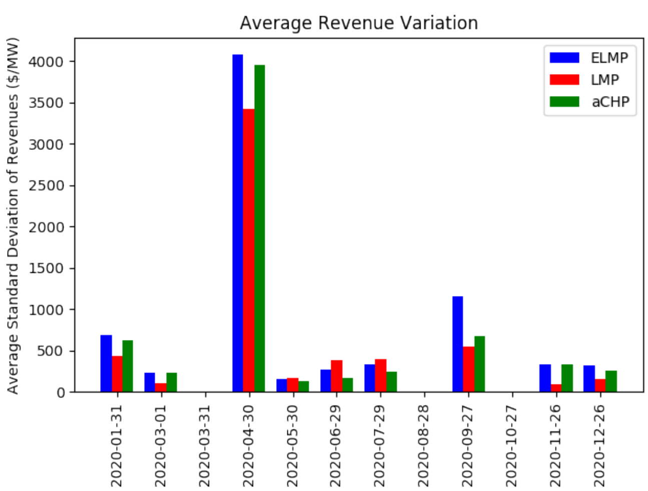

Fig. 9. Averages of the standard deviations of the revenues across near-optimal solutions in the RTS GMLC Network power system according to various pricing schemes.

by one pricing scheme valuing these two generators more than another pricing scheme.

In comparing the no-network case with the network case, we broadly observe that on average, the network case exhibits lower variation, but not by a significant margin. This coincides with our observation that the network case in general offers fewer solutions and hence we would expect less variation. We would need to observe more cases to make this claim with more certainty however.

IV. Conclusion

The presence of both small and symmetric (or almost-symmetric) generators of a power system introduces the potential for there to be a large, diverse set of near-optimal solutions to the UC problem. When analyzed from an economic perspective, near-optimal solutions can exhibit great disparity in how they distribute revenues and profits among generators. Especially in the case of power systems with a large number of generators with a high proportion of low-power generators (as in the CAISO system), there may exist solutions within an optimality gap of 0.01% where revenue and profits may vary dramatically.

Minimizing the volatility of revenues and profits among generators can be brought about in part by reducing the size of the MIP gap so as to reduce the number of solutions (to 1 if possible) which ideally will eliminate highly disparate solutions. Alternatively, certain pricing schemes can exhibit lower levels of variation. Empirically, we have seen that in the CAISO system, the aCHP pricing scheme expresses the least variation, in part due to its solution-independent prices. The choice of pricing scheme may reduce the presence of significant variation.

References

[1] Gurobi Optimization, Inc., "Gurobi optimizer reference manual," 2018. [Online]. Available: http://www.gurobi.com

[2] J. Ostrowski, M. F. Anjos, and A. Vannelli, "Modified orbital branching for structured symmetry with an application to unit commitment," *Mathematical Programming*, vol. 150, no. 1, pp. 99–129, 2015.

[3] B. Knueven, J. Ostrowski, and J.-P. Watson, "Exploiting identical generators in unit commitment," *IEEE Transactions on Power Systems*, vol. 33, no. 4, 2018.

[4] P. Bendotti, P. Fouilhoux, and C. Rottner, "Sub-symmetry-breaking inequalities and application to the unit commitment problem," *HAL*, vol. 2019, 2019.

[5] B. Knueven, J. Ostrowski, and J.-P. Watson, "On mixed-integer programming formulations for the unit commitment problem," pp. 857–876, 2020.

[6] M. Fischetti and L. Liberti, "Orbital shrinking," in *International Symposium on Combinatorial Optimization*. Springer, 2012, pp. 48–58.

[7] B. Eldridge, R. O'Neill, and B. Hobbs, "Pricing in day-ahead electricity markets with near-optimal unit commitment," 2018.

[8] D. A. Schiro, T. Zheng, F. Zhao, and E. Litvinov, "Convex hull pricing in electricity markets: Formulation, analysis, and implementation challenges," *IEEE Transactions on Power Systems*, vol. 31, no. 5, pp. 4068–4075, 2016.

[9] B. Hua and R. Baldick, "A convex primal formulation for convex hull pricing," *IEEE Transactions on Power Systems*, vol. 32, no. 5, pp. 3814–3823, 2017.

[10] C. Wang, P. B. Luh, P. Gribik, L. Zhang, and T. Peng, "The subgradient-simplex based cutting plane method for convex hull pricing," in *IEEE PES General Meeting*. IEEE, 2010, pp. 1–8.

[11] P. R. Gribik, W. W. Hogan, and S. L. Pope, "Market-clearing electricity prices and energy uplift," *Cambridge, MA*, 2007.

[12] B. Knueven, J. Ostrowski, and J. Wang, "The ramping polytope and cut generation for the unit commitment problem," *INFORMS Journal on Computing*, vol. 30, no. 4, pp. 739–749, 2018.

[13] Y. Guan, K. Pan, and K. Zhou, "Polynomial time algorithms and extended formulations for unit commitment problems," *IISE transactions*, vol. 50, no. 8, pp. 735–751, 2018.