LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-599856

Large-scale Stochastic Optimization for Unit Commitment and Economic Dispatch

Y. Yao, C. Meyers, S. Smith, T. Edmunds

November 8, 2012

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Large-scale Stochastic Optimization for Unit Commitment and Economic Dispatch

Introduction

Our FY 2012 LDRD research on stochastic optimization with its application in power grid planning and operations has advanced the state-of-the-art of this area in the following aspects:

(1) Deploying stochastic unit commitment on large real-world applications,
(2) Stochastic optimization coupled with real-time production cost simulation (interleaved), and
(3) Real-time stochastic optimization with rolling-horizon look-ahead.

In conjunction, a number of computational experiments have been performed to evaluate the capability of the latest stochastic optimization software and explore opportunities for future research.

In what follows, we discuss each of these contribution areas, including the results of our computational experiments and their implications on the scalability and solvability of large-scale stochastic optimization models.

Problem Description

The base power grid model that we use is one developed by the California Independent Systems Operator (CAISO) as part of a study they performed on operational requirements and market impacts at a 33% renewable portfolio standard (RPS). The state of California has legislated transitioning to a 33% RPS by the year 2020[1], which means that fully 1/3 of the state's electricity must come from renewable sources. The CAISO study seeks to quantify the impacts of this mandate and identify operational requirements needed to run the state energy grid under this standard.

The CAISO power grid model is a large-scale model of the Western Electricity Coordinating Council (WECC) power grid, which covers the entire western United States. The WECC model consists of 2486 generators (including renewable energy sources) attached to 46 nodes representing transmission hubs (see Figure 1). These nodes are connected with 120 transmission lines, forming a zonal-level grid. The CAISO model includes definitions of reserves for ancillary power, including load following, regulation, spinning, and non-spinning reserve requirements. It also includes 7 storage facilities, with head and tail storage. Finally, the model contains numerous specialized classes of constraints tuned to the real-world operation of the California power grid.

[1] Bill SBX1 2, signed by Governor Brown in 2011; http://gov.ca.gov/news.php?id=16974.

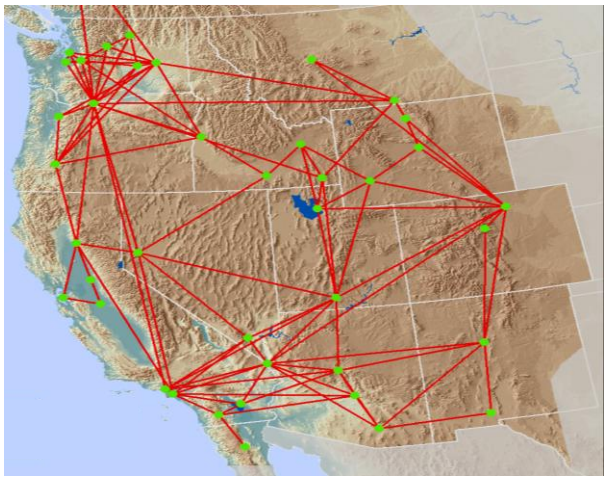

Figure 1. Model of the WECC used in the CAISO 33% RPS study. Nodes are represented by green dots and transmission lines by red lines.

To assess and compare the impact of different types of future renewable energy buildouts, the CAISO model contains 7 different renewable energy profiles. Load data is based on realized loads from the years 2005 and 2006, scaled to the year 2020[2]. Reserve requirements were calculated via a separate simulation analysis.

The mathematical optimization problem that the CAISO model was deigned to answer is a *unit commitment and dispatch* problem: determine which generators need to be turned on, at which times, to satisfy the forecasted demand requirements across the grid, obeying storage and transmission capacities, and satisfying reserve requirements and side constraints. Conventionally this problem can be decomposed into two strongly coupled decisions:

- unit commitment: which generators should be turned on, and at what times?
- dispatch: given the unit commitment decisions of each generator, how should power be allocated to satisfy demand across the grid?

In practice, CAISO and other ISOs solve these unit commitment and dispatch problems using a mixed-integer programming approach, in which unit commitment and dispatch values are treated as variables, the system operating conditions are encoded as constraints, and the objective is to minimize system operating cost.[3]

---

[2] California Independent System Operator: Integration of Renewable Resources – Operational Requirements and Generation Fleet Capability at 20% RPS, 2010; http://www.caiso.com/2804/2804d036401f0.pdf. Although the modeling methodology in the 20% RPS study is significantly different, the underlying data is substantially similar.
[3] For a full mathematical programming formulation of this problem, see, for instance, Padhy: Unit Commitment – A Bibliographic Survey. *IEEE Transactions on Power Systems*, 19(2):1196-1205, 2004, http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=01295033.

Stochastic Optimization

Traditionally, ISO unit commitment decisions are made using a reserves adjustment method, in which uncertainties in load and renewable generation are accounted for by imposing extremely conservative reserve requirements. This causes the grid to substantially overcommit resources as a means of ensuring that there is nearly always enough generation capacity to meet realized demand. More recently, the approach of *stochastic optimization* has been employed as a means of more accurately accounting for these uncertainties and lessening the problem of overcommitment. In a stochastic unit commitment approach, each of the uncertain parameters is associated with a set of possible scenarios encompassing the range of its possible outcomes. Each of these scenarios is paired with a probability that represents the likelihood that such a scenario will be realized. A bi-level mixed-integer program is used to determine the set of least-cost unit commitments that probabilistically account for each of the possible scenario realizations.

Deploying stochastic unit commitment on large real-world applications

In our research, we have adapted the CAISO 33% RPS model to solve for stochastic unit commitments, improving upon the reserves adjustment approach used in the original study. This involved two major components:

- developing a set of scenarios to capture the uncertainties in load and renewable generation; and
- implementing a stochastic unit commitment algorithm over these scenarios.

For the first step, we leveraged the work of an ongoing study for the California Energy Commission[4] on evaluating demand response and storage alternatives. This study employs the same underlying CAISO 33% RPS model. A novel feature of this study is that it uses an atmospheric model to determine possible weather trajectories, which are used to derive day-ahead load and solar/wind generation forecasts. The atmospheric forecasts are generated using the Weather Research and Forecast (WRF) model, which is a fully compressible multi-physics atmospheric modeling software package. A statistical sampling method is used to convert a large suite of forecasts into a smaller set of weather scenarios, each associated with a (potentially different) probability. These weather scenarios are then used to determine the day-ahead load (via a dL/dTemp transformation), solar (from shortwave downward radiative flux), and wind (from wind speed/power curves) generation forecasts.

To implement the stochastic unit commitment algorithm, we employ some novel capabilities of the PLEXOS software package. PLEXOS[5] is a power market simulation package based on mathematical programming that is capable of modeling many different areas of the energy sector at varying time

[4] Edmunds, T., Lamont, A., Schmidt, A., Mirocha, J., Simpson, M., Top, P., Smith, S., and Bulaevskaya, V.: Probabilistic Renewable Generation Scenarios for Analysis of Demand Response and Storage Resources. California Energy Commission contract 500-10-051, 2012.
[5] www.energyexemplar.com

horizons; this package was used for production simulation runs in the CAISO 33% RPS analyses. We worked with the PLEXOS code developers to extend the underlying mathematical programming formulation to accommodate stochastic unit commitment over our set of load and solar/wind generation scenarios. The resulting bi-level mixed-integer program treats the commitments of slow-start generators as first-stage decisions (common to all scenarios), and the commitments of fast-start generators and economic dispatch as second-stage decisions (potentially different between scenarios).

We next performed a number of experiments to address the computational demands associated with running large-scale stochastic unit commitment problems. These experiments focused on two quantities of interest:

- the relationship between the *tolerance level* (degree of optimality) required by the underlying mixed-integer programming algorithm, and solution times; and
- the relationship between the *number of forecast scenarios* and solution times.

Both of these relationships are key to the real-world applicability of stochastic unit commitment algorithms.

Optimality Tolerance

We evaluated the impact of changing the mixed-integer programming optimality tolerance on the quality of the resulting optimal solution, as shown in Table 1. In this table, the operational cost column measures the quality of the optimal solution obtained. The 0.05% tolerance level is used as a default at CAISO, which is in accordance with the tolerance level established for their daily operational runs.

Table 1. Summary of optimality tolerance experiments, based on two scenarios for July 21, 2020.

Tolerance (%)	Time (min)	Cost	Delta Cost %
0.05	13.8	88042.94	0
0.1	13.75	88042.94	0
0.25	10.1	88246.26	0.23
0.5	9.9	88465.73	0.48

We observe that even if the tolerance is raised by an order of magnitude (from 0.05% to 0.5%), the overall cost increases by less than 0.5%. *We conclude that a larger (say, 0.5%) tolerance level can feasibly be employed to reduce computation time without a significant deterioration in performance, which is especially important for runs with many scenarios*. This relationship can clearly be seen in Figure 2, which is based on the data from Table 1. Here we observe that the execution time decreases by around 30% as the tolerance is varied from .05 to .5, with a negligible increase in overall cost.

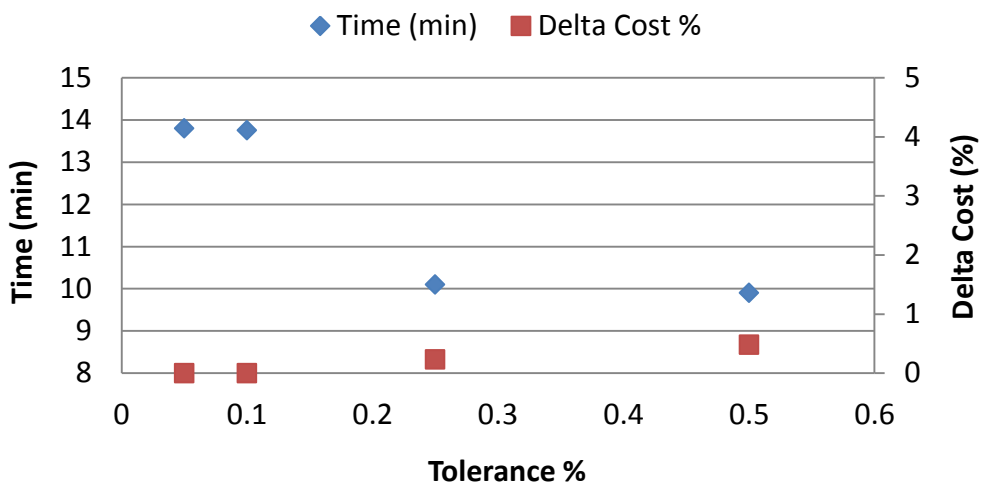

Figure 2. Graph of optimality tolerance experiments, based on two scenarios for July 21, 2020.

These findings are particularly relevant given the large amount of uncertainty in the underlying data for the CAISO 33% RPS model, much of which is based on projections for the year 2020. The tolerance value of .05% may be appropriate for daily operational models in which CAISO is accountable for even a small percentage of error, but we argue that planning models can be feasibly run at lower tolerance levels with negligible difference, especially given the inherent uncertainty of such results. These findings are in agreement with previous observations on the optimality tolerance associated with the non-stochastic version of the CAISO 33% RPS model.[6]

Number of Forecast Scenarios

Table 2 summarizes our computational findings on the relationship between the number of scenarios, solution time, and peak memory usage.

Table 2. Summary of scenario number experiments. Based on two scenarios for July 21, 2020 and optimality tolerance = .05%. For the 12-scenario case, the run was terminated with a feasible solution within 4.8% of optimality.

Number of Scenarios	Solution Time (min)	Peak Memory Usage (GB)
2	13.8	1.7
4	106	5.8
6	207	7.5
8	268	8
10	331	9.8
12	456	13.2

[6] Yao, Y., Meyers, C., Schmidt, A., Smith, S., and Streitz, F.: Observations on the Optimality Tolerance in the CAISO 33% RPS Model. Technical Report LLNL-TR-501034, Lawrence Livermore National Laboratory, 2011.

From the table, we observe that PLEXOS can solve stochastic unit commitment problems with 10 scenarios or fewer within a 6-hour time window. Determining commitments for 10 scenarios for a monthly timespan would take about 7 days, and for a year would take 84 days, assuming these run times are representative of the average. *This strongly suggests that the number of scenarios used in stochastic optimization should be limited to a number for which the solution time is not considered an excessive burden*.

Finally, we note that the execution times of daily runs can vary significantly over time, as demonstrated in Table 3.

Table 3. Summary of daily experiments, based on four scenarios and optimality tolerance = 0.05%.

Date	Time (min)
7/2/2020	32
7/3/2020	25
7/4/2020	76
7/5/2020	34
7/6/2020	135
7/7/2020	135
7/8/2020	70

Here we observe that there can be a factor of four separating the solution times for consecutive days. While this finding may seem troublesome, it is not entirely unexpected in the mixed-integer programming domain, in which small changes in problem formulations can lead to large differences in solution times.[7] Fortunately, subsequent experimentation has shown that the differences are not usually much greater than those observed here, and the times in Table 2 are indeed representative of a fairly 'average' day. All of these factors must be kept in mind when deploying stochastic optimization methods on large-scale real-world problems.

Stochastic Optimization Coupled with Real-Time Production Cost Simulation (Interleaved)

In the real-world CAISO power grid, unit commitment and dispatch decisions are naturally made on differing time scales: unit commitments of slow-start generators are optimized and set the day before operation, while dispatch and fast-start generation are determined and modified throughout the course of the day, in response to demand and the realization of uncertainties in renewable generation. Most grid optimization models fail to realistically capture this time-span dichotomy, or consider only part of the problem (such as optimizing over dispatch while assuming unit commitments are fixed). In our research we have identified a method for optimizing over both time spans simultaneously, filling a gap in the grid optimization literature.

[7] See, for instance, Danna, E.: Performance Variability in Mixed-Integer Programming. In *5th Workshop on Mixed-Integer Programming*, 2008, http://coral.ie.lehigh.edu/~jeff/mip-2008/talks/danna.pdf.

Our method is as follows.

1) First, we produce an optimal day-ahead *hourly* (60-minute) unit commitment schedule for slow-start generators over all forecast scenarios;
2) Next, we perform *intra-hour* (5-minute) unit commitment on fast-start generators and economic dispatch (ED) based on the realized scenario.

For each day, the hourly unit commitment schedule from the unit commitment run dictates the commitment status of the long-start generators in the subsequent economic dispatch run. Within each hour, additional runs are made at 5-minute timesteps to determine the dispatch values and the commitment status of fast-start generators. This is reasonably close to the way that CAISO actually operates, minus intra-minute grid stability analyses that are not well captured by a mixed-integer programming approach.

Performing sequential unit commitment/economic dispatch daily runs require that the system status at the end of a day is passed as initial conditions to the next day's optimization decisions. This is necessary to avoid violations in generator ramp rates in transitioning from one day to the next. We employ data stitching tools to pass the unit commitment schedule and system status from one day to the next and to set up the corresponding PLEXOS runs correctly. The PLEXOS file setup for this two-stage optimization approach is outlined in Figure 3.

- For all generators, extract the values for Initial Generation, Initial Units Generating, Initial Hours Up, and Initial Hours Down from a file. Set all values in these "initial state" files to 0 at the beginning of day 1.
- In Model->Report check the box next to "Text Files", and select Generation, Unit Generating, Hours Up, and Hours Down to output in every period in "state output" files.
- Set load and wind/solar input files to "forecast" data in the "stochastic unit commitment" tag.
- Set Generator's Commitment to the Unit Generating state output file in the "15-minute dispatch" tag.
- Set load and wind/solar input files to "realized" data in the "15-minute dispatch" tag.

Figure 3. Specification of PLEXOS file setup for two-stage interleaved optimization.

The optimization setup for running PLEXOS in sequence with stitching tools to perform daily stochastic unit commitment and sub-hourly economic dispatch is outlined in Figure 4.

For each day in the planning horizon:

- Run the prescheduling (PASA) and medium-term (MT) model to set constraints for the short-term (ST) model;
- Run the ST model to obtain stochastic hourly unit commitment decisions;
- Run the intra-day stitching process to pass generator commitments to the economic dispatch model;
- Run the intra-hour economic dispatch model at a 5-minute time step; and
- Run inter-day stitching to pass the system's ending status to the start of the next day.

Figure 4. Specification of PLEXOS model sequencing for interleaving unit commitment and dispatch.

With these specifications, it is possible in PLEXOS to run a simulation of the CAISO grid at a greater degree of realism than was previously achievable. We have worked extensively with the PLEXOS developers to vet and test this solution process and diagnosed and fixed many problems along the way.

Real-Time Stochastic Optimization with Rolling Horizon Look-ahead

To further improve the realism and solution quality produced by the CAISO model, we assessed the benefits of adding a *look-ahead* component to the solution process. In a look-ahead strategy, the solution at one time interval is generated using partial information about the next time period. A real-world analog might involve using forecasted demands from day 2 to influence the end-of-day ramping behavior of generators on day 1. These kinds of strategies are regularly employed by CAISO to help improve overall system cost.

Mathematically, the mixed-integer program from the first time period is augmented with a linear programming relaxation for each of the look-ahead time periods. Including information about the subsequent time periods allows the solver to find a potentially better solution overall, while limiting the augmented periods to linear programming relaxations helps mitigate the computational demands associated with this look-ahead. In subsequent time periods, the horizon 'rolls' forward, so that the mixed-integer program for the second time period is solved using the same number of look-ahead periods as the first, and thus including the data for one additional time period.

In our analyses, we considered a look-ahead on two time horizons:

- Day-ahead hourly stochastic unit commitments of slow-start generators, with look-ahead on a daily timestep; and
- 5-minute unit commitment of fast-start generators and economic dispatch, with look-ahead on a 5-minute timestep.

Table 4 gives our results for runtimes for look-ahead of various lengths applied to day-ahead hourly stochastic unit commitments of slow-start generators, using 6 stochastic wind samples.

Table 4. Runtimes of stochastic unit commitment for various look-ahead lengths. Stochastic samples are based on six wind scenarios in the Pacific Gas & Electric and Southern CA Edison service regions; optimality tolerance = 0.5%.

	Real Time UC/ED Look-ahead Periods				
	0	1	2	5	Mean
15-Jan-20	0:56:14	0:55:33	1:05:35	1:30:06	1:06:52
29-Apr-20	2:37:01	2:30:41	2:42:46	1:36:34	2:21:45
21-Jul-20	1:34:29	1:41:19	1:49:05	2:40:36	1:56:22
1-Oct-20	1:13:52	1:21:21	1:26:50	1:29:50	1:22:58
Mean	1:35:24	1:37:14	1:46:04	1:49:16	1:42:00

Here, we observe that performing look-ahead for a moderate number of time periods (1 or 2) does not increase the computational demands significantly. Extending to a look-ahead of 5 time periods can increase the solution time by up to a factor of 50%, but this is not consistent and a look-ahead of this length should also be considered potentially practical.

Table 5 gives our results for total system costs for 5-minute unit commitment of fast-start generators and economic dispatch based on the realized scenario, with look-ahead of varying time steps.

Table 5. Total system costs of day-ahead stochastic UC and 5-minute UC/ED for various look-ahead lengths. Total costs include all generators' start-up, generation, and emission costs.

	Real Time UC/ED Look-ahead Periods				
	0	1	2	5	Mean
15-Jan-20	9.2703E+04	9.2848E+04	9.2821E+04	9.2815E+04	9.28E+04
29-Apr-20	7.0849E+04	7.0537E+04	7.0388E+04	7.0038E+04	7.05E+04
21-Jul-20	1.2360E+05	1.2384E+05	1.2381E+05	1.2390E+05	1.24E+05
1-Oct-20	9.8403E+04	9.8586E+04	9.8656E+04	9.8421E+04	9.85E+04
Mean	9.6388E+04	9.6452E+04	9.6419E+04	9.6294E+04	9.64E+04

Surprisingly, we do not observe a significant cost reduction by adding look-ahead (of any length), which is counterintuitive. Our investigation suggests that the ineffectiveness of this look-ahead may be due to the small intra-hour variability of the synthetic load data used in the experiments. One of the main advantages of look-ahead that it allows preemptive adjustment of ramp rates and generator levels in anticipation of future fluctuations in demand, but without these fluctuations, the potential benefit is likely to be much smaller. Thus, it seems our dataset is not well-suited to take advantages of the benefits that look-ahead can afford; however, this may be a valuable technique for other datasets and should still be considered a potentially valuable tool to add to the modeling arsenal.

Conclusions and Future Research

We have described our progress in applying stochastic optimization techniques to real-world power grid planning problems, advancing the state-of-the-art in terms of both application size and modeling methodology. In the future, we will be deploying these techniques on a full-scale yearly optimization of CAISO grid operations, which will undoubtedly lead to further avenues of research. We also have an active engagement with IBM and we anticipate solving these problems on a supercomputing platform using a distributed-memory solver currently under development. This may require further decomposition of the underlying mixed-integer programming formulation in order to effectively leverage the work assigned to each of the supercomputing nodes. We expect the products of this research to find immediate application in ongoing engagements with the California Energy Commission and utilities in the state of California.