Applying High Performance Computing to Transmission-Constrained Stochastic Unit Commitment for Renewable Energy Integration

Anthony Papavasiliou, *Member, IEEE,* Shmuel S. Oren, *Fellow, IEEE,* and Barry Rountree

Abstract—We present a parallel implementation of Lagrangian relaxation for solving stochastic unit commitment subject to uncertainty in renewable power supply and generator and transmission line failures. We describe a scenario selection algorithm inspired by importance sampling in order to formulate the stochastic unit commitment problem and validate its performance by comparing it to a stochastic formulation with a very large number of scenarios, that we are able to solve through parallelization. We examine the impact of narrowing the duality gap on the performance of stochastic unit commitment and compare it to the impact of increasing the number of scenarios in the model. We report results on the running time of the model and discuss the applicability of the method in an operational setting.

Index Terms—Stochastic optimization, unit commitment, Lagrangian relaxation, scenario selection, parallel computing.

I. INTRODUCTION

The large-scale integration of renewable energy sources and demand response resources in power systems is increasing the complexity of power systems planning and operations. This is due to the high level of uncertainty, as well as the increased number of agents participating in the active coordination of supply and demand. An important technical problem that arises as a result of this paradigm shift in power system operations is the requirement for systematic methods of committing reserves in order to ensure the reliable operation of the system. The commitment of reserves is currently performed in an ad hoc fashion [1]. Methodical approaches for committing reserves based on stochastic optimization, robust optimization, security constraints and probabilistic constraints have been proposed recently. The advantage of these approaches relative to current practice is the representation of uncertainty within the optimization model which, however, results in optimization problems that are more challenging computationally. The solution of these problems within operationally acceptable time frames can strongly benefit market clearing and dispatch methods. The computational methods proposed for addressing these problems often rely on decomposition methods, and can therefore be distributed across multiple processor cores.

The purpose of this research is to perform a scoping study for short-term scheduling under uncertainty in anticipation of

A. Papavasiliou is with CORE and the department of Mathematical Engineering at the Catholic University of Louvain, Louvain la Neuve, 1348, Belgium (e-mail: anthony.papavasiliou@uclouvain.be).

S. S. Oren is with the Department of Industrial Engineering and Operations Research, UC Berkeley, CA, 94720 USA (oren@ieor.berkeley.edu).

B. Rountree is with the Lawrence Livermore National Laboratory, Livermore, CA, 94550 USA (e-mail: rountree4@llnl.gov).

the technological progress that is expected to take place in the coming years in the area of parallel computation. Our objective in the current publication is to identify the appropriate scale for stochastic unit commitment formulations by analyzing the sensitivity and performance of the resulting unit commitment policies to the number of scenarios considered in the model.

A. Literature Review

Various modeling approaches have been recently set forth in order to deal with power system operations in a systematic fashion. Stochastic unit commitment models [2], [3] represent uncertainty in terms of appropriately weighted scenarios, where the objective of the unit commitment problem is to minimize expected cost of operations. The resulting models are large scale and require a substantial amount of information regarding the source of uncertainty in the model. These models are typically solved through Lagrangian relaxation [4], [5], augmented Lagrangian methods [3] or progressive hedging methods [2], [6] and, consequently, the parallelization of the solution algorithms is straightforward. Robust optimization models [7] strive to optimize the cost of operations against the worst possible realization of uncertainty within a prescribed set. A setback of robust unit commitment models is that the resulting solution algorithms which are based on Benders' decomposition result in bilevel optimization problems and their parallelization is not straightforward. Various models adopt security criteria [8], [9], where the system is expected to withstand the failure of system components without shedding load. These models are addressed via Lagrangian relaxation and Benders decomposition. Probabilistic constraints [10] have also been proposed for the unit commitment problem. Despite the recent advancements of mixed integer linear programming algorithms that has led to their adoption in market clearing and system operations, these algorithms cannot be applied directly to stochastic unit commitment due to memory and running time limitations but are instead employed as modules in the aforementioned decomposition algorithms.

a) Lagrangian Relaxation and Scenario Selection for Stochastic Unit Commitment: Among the methods previously described for systematically addressing short-term scheduling uncertainty, the focus of this paper is on stochastic programming. Stochastic formulations of the unit commitment problem typically relax the non-anticipativity constrains. The first Lagrangian decomposition algorithm for stochastic unit commitment was introduced by Takriti and Birge [2] who use

the progressive hedging algorithm for solving the problem. The authors in [2] relax the optimization problem at each node of the scenario tree. Instead, Carpentier et al. [3] relax the market clearing constraint and solve a separate stochastic program for each unit in the system. A similar relaxation is adopted by Nowak and Römisch [5]. In this paper we use a scenario decomposition approach.

The performance of the stochastic unit commitment model depends critically on the set of scenarios that are input in the formulation and how they are weighed relative to each other. Previous research on scenario selection exploits results on the distances among probability measures to cast the scenario selection problem as a transportation problem [11], [12] with the purpose of minimizing the distance among the probability measures of the full and reduced problem respectively. These scenario selection algorithms have been applied to hydrothermal scheduling [13]. However, these algorithms are only applicable for continuous sources of uncertainty, rather than component failures, and also ignore the objective function of the problem that is being optimized, which in our case leads to poor performance due to the great sensitivity of the objective function to load shedding. A detailed discussion and computational study of the performance of these algorithms in stochastic unit commitment is presented in previous work by the authors [14], [15]. Moment matching methods [16], [17] suffer from the same drawbacks. By exploiting distributed computation we can address problems with a large number of samples. For this reason we implement the sample average approximation (SAA) algorithm, introduced by Kleywegt et al. [18], which has asymptotic guarantees of convergence to optimality for large sample sizes. The SAA algorithm studied by Kleywegt et al. [18] addresses two-stage stochastic programs with discrete first-stage decisions and therefore fits the description of our model. We test the performance of the SAA algorithm against a scenario selection algorithm inspired by importance sampling that was introduced in [14] and is briefly presented in this paper.

b) Distributed Computing: Distributed computation has a rich history in the area of power systems planning and operations. One early application is presented by Monticelli et al. [19] who parallelize Benders decomposition for solving the security constrained optimal power flow problem with corrective rescheduling. Pereira et al. [20] exploit distributed computing for the reliability evaluation of composite outages, for scenario analysis in a hydro dominated system and for the application of Benders decomposition in security constrained dispatch. Kim and Baldick [21] develop a parallel augmented Lagrangian relaxation algorithm for solving distributed optimal power flow across multiple operating regions. Bakirtzis and Biskas [22] extend the work of Kim and Baldick [21] by proposing an alternative Lagrangian relaxation for the optimal power flow problem and develop a parallel implementation of the algorithm [23]. A review paper for the application of high performance computing in power systems planning and operations is presented by Falcao [24].

B. Paper Contributions

The existing work on parallel computing for power systems planning and operations is focused on simulation, static instances of the optimal power flow problem and the midterm (monthly) scheduling of hydro-thermal production. The recent proliferation of renewable energy sources such as wind and solar power has stimulated interest in dealing with the short-term (hourly) uncertainty caused by these resources. Although numerous researchers have studied the integration of renewable power supply on short-term operations [25], [26], [15], none have exploited high performance computing for the short-term scheduling problem. The use of high performance computing enables us to study instances of the problem the size of which has not been addressed before. This enables us to address questions regarding the influence of the size of the scenario set on the unit commitment policy, the bounds of the problem and the performance of the resulting policy.

The decomposition algorithm and relaxation used in this paper have been developed in previous work by the authors [15], [14]. The focus of the authors' previous work was the validation of the proposed scenario selection model against operating practice for reserve commitment based on deterministic unit commitment in order to develop a consistent framework for quantifying reserves in systems with significant levels of renewable energy and demand response penetration. The present paper is focused, instead, on computational aspects of stochastic unit commitment and the applicability of the model in an operational setting. Parallel computing is an enabling technology in the adoption of stochastic optimization and other systematic approaches (e.g. robust optimization, probability constraints and security constraints) in short-term power systems scheduling. The purpose of this paper is to demonstrate this by solving a model the scale of which is an order of magnitude larger than the state of the art in stochastic unit commitment. We also investigate practical tradeoffs between running time and the modeling of uncertainty.

This work addresses the following specific questions: (a) How does the performance of our scenario selection algorithm scale with the number of scenarios? How does our scenario selection algorithm compare relative to the sample average approximation (SAA) algorithm of Kleywegt et al. [18]? (b) How does the performance of stochastic unit commitment depend on the duality gap of the non-anticipativity constraints? What is the relative importance of duality gaps and scenario size on the performance of the algorithm? (c) Can distributed computing scale the running time of our model down to a level where it becomes acceptable for operational purposes?

The remaining paper is organized as follows. In Section II we formulate the two-stage stochastic unit commitment model. We then present a parallel implementation of Lagrangian relaxation for solving the problem in Section III, and a scenario selection algorithm for formulating the model in Section IV. We conclude our analysis with a study of a reduced model of the Western Electricity Coordinating Council in Section V and report our computational experience.

II. Model Description

The focus of this paper is the day-ahead scheduling of generators subject to real-time renewable power supply uncertainty and outages of transmission lines and generators. The problem is cast as a two-stage optimization, where the first stage represents day-ahead decisions and the second stage represents the real-time recourse to the revealed system conditions.

In the following model formulation, u represents a binary variable indicating the on-off status of a generator, v is a binary startup variable and p is the production level of each generator. The minimum load cost of a generator is denoted as K_g, the startup cost as S_g and the constant marginal cost as C_g. The set of generators is denoted as G. The model that we present in this paper accounts for transmission constraints, with power flows over transmission lines denoted as e and bus angles denoted as θ. L represents the set of transmission lines. The demand for each hour t at each bus of the network n is denoted as D_{nt}, where N is the set of buses in the network. Operating constrains are denoted compactly in terms of a feasible set $\mathcal{D}$, and vectors are denoted in bold. Thus, the notation $(\mathbf{p}, \mathbf{e}, \theta, \mathbf{u}, \mathbf{v}) \in \mathcal{D}$ encapsulates the minimum/maximum run limits, minimum up/down times and ramping rate limits of generators, as well as Kirchhoff's voltage and current laws and the thermal limits of lines. The constraint set $\mathcal{D}$ can also include reserve requirements and flow constraints that protect the network against load and renewable supply forecast errors, as well as generator and transmission line outages. We use a DC approximation of the voltage and current laws, which is a standard assumption in the unit commitment literature [8] and is also consistent with the day-ahead unit commitment model used by the system operator.

The objective is to minimize the cost of serving forecast demand. The problem in the deterministic setting (assuming an accurate forecast of renewable power production and component failures) can be described as follows:

$$(UC): \min \sum_{g \in G} \sum_{t \in T} (K_g u_{gt} + S_g v_{gt} + C_g p_{gt}) \quad (1)$$

s.t.
$$\sum_{g \in G_n} p_{gt} = D_{nt}, n \in N, t \in T \quad (2)$$

$$P_g^- u_{gt} \leq p_{gt} \leq P_g^+ u_{gt}, g \in G, t \in T \quad (3)$$

$$e_{lt} = B_l(\theta_{nt} - \theta_{mt}), l = (m, n) \in L, t \in T \quad (4)$$

$$(\mathbf{p}, \mathbf{e}, \theta, \mathbf{u}, \mathbf{v}) \in \mathcal{D}, \quad (5)$$

where B_l represents the susceptance of line l and P_g^-, P_g^+ represent the minimum and maximum run limits of generator g. The set of generators located in each bus n is denoted by G_n.

The stochastic formulation involves a two-stage process, where the set of uncertain outcomes is represented as S. First-stage unit commitment and startup decisions are represented respectively as w and z and are defined for slow-responding units G_s for which commitment decisions need to be made

in advance, in the day-ahead time frame. The problem to be solved is the following:

$$(SUC):$$
$$\min \sum_{g \in G} \sum_{s \in S} \sum_{t \in T} \pi_s (K_g u_{gst} + S_g v_{gst} + C_g p_{gst}) \quad (6)$$

s.t.
$$\sum_{g \in G_n} p_{gst} = D_{nst}, n \in N, s \in S, t \in T \quad (7)$$

$$P_{gs}^- u_{gst} \leq p_{gst} \leq P_{gs}^+ u_{gst}, g \in G, s \in S, t \in T \quad (8)$$

$$e_{lst} = B_{ls}(\theta_{nst} - \theta_{mst}),$$
$$l = (m, n) \in L, s \in S, t \in T \quad (9)$$

$$u_{gst} = w_{gt}, v_{gst} = z_{gt}, g \in G_s, s \in S, t \in T \quad (10)$$

$$(\mathbf{p}, \mathbf{e}, \theta, \mathbf{u}, \mathbf{v}) \in \mathcal{D}_s, \quad (11)$$

where decision variables are now contingent on the scenario $s \in S$. Note that the domain $\mathcal{D} = \times_{s \in S} \mathcal{D}_s$ is decomposable across scenarios. Each scenario is weighed in the objective function by a probability π_s. The sources of uncertainty in the model include: (a) The hourly supply of renewable energy, which is reflected in the net demand D_{nst}. (b) The availability of a generator g. If a generator fails for scenario s, then $P_{gs}^- = P_{gs}^+ = 0$ (assuming the outage lasts for the entire horizon). (c) The availability of a line l. If a lines fails for scenario s, then $B_{ls} = 0$ (assuming the outage lasts for the entire horizon and the line is removed from service by circuit breakers). The selection and weighing of a representative set of scenarios is explained in Section IV. The generation of spatially correlated renewable supply scenarios, generation and transmission line outages is described in Section V.

Upon realization of the uncertain outcome, fast generators G_f adjust their unit commitment schedule u_{gst}, whereas slow generators G_s are forced to maintain their day-ahead unit commitment schedule, as indicated by the non-anticipativity constraints of Eq. (10). All generators can adjust their production levels according to the realization s, regardless of whether they are slow-responding or fast-responding resources. This two-stage stochastic unit commitment formulation follows the model of Ruiz et al. [26].

The fact that transmission constraints are accounted for in Eq. (9) and in Eq. (11) through the flow limits on lines, $-TC_l \leq e_{lst} \leq TC_l$ implies that the model is appropriate for optimizing the placement of reserves in the network. Previous work [14] has demonstrated that transmission constraints are crucial in properly quantifying reserve requirements and operating costs resulting from large-scale renewable energy integration. The fact that the stochastic unit commitment model accounts for transmission constrains necessitates a time series model that accounts for spatial correlations of renewable production, and increases the computational challenges of solving the model in terms of scenario selection and decomposition. These challenges have been addressed in previous work [14]. The focus of the present paper is to exploit parallel computing in order to assess the sensitivity of the model on the number of scenarios.

III. Solution Methodology

In this section we present a parallel algorithm that has been implemented on the High Performance Computing facility at the Lawrence Livermore National Laboratory in order to study the tradeoff of duality gaps and scenario set size in stochastic unit commitment and examine the benefit of parallel computing. The results of Section V-B indicate that alternative relaxations and dual function optimization methods can yield superior performance, which reinforces the argument that parallel computing holds great promise for power system operations under uncertainty. A detailed analysis and comparison of alternative relaxations and dual function optimization schemes is deferred for future research.

The non-anticipativity constraints of Eq. (10) couple decisions across scenarios. In particular, the commitment decisions u_{gst} among different scenarios for slow units G_s have to be consistent with the day-ahead commitment decisions w_{gt} for slow units, and the same holds true for startup decisions. The Lagrangian relaxation algorithm relies on the observation that the relaxation of the non-anticipativity constraints in Eq. (10) results in unit commitment subproblems (UC_s), as in Eqs. (1) - (5), that are independent across scenarios. The Lagrangian dual function is obtained as:

$$\mathcal{L} = \sum_{s \in S} \pi_s (\sum_{g \in G} \sum_{t \in T} (K_g u_{gst} + S_g v_{gst} + C_g p_{gst})$$
$$+ \sum_{g \in G_s} \sum_{t \in T} (\mu_{gst}(u_{gst} - w_{gt}) + \nu_{gst}(v_{gst} - z_{gt})) \quad (12)$$

The problem is solved by maximizing the Lagrangian dual function. The solution of the Lagrangian involves one second-stage unit commitment problem for each scenario (UC_s), and one first-stage optimization $(P1)$. The first-stage optimization is formulated as:

$$(P1): \max \sum_{g \in G_s} \sum_{s \in S} \sum_{t \in T} (\mu_{gst} w_{gt} + \nu_{gst} z_{gt}) \quad (13)$$
$$\text{s.t.}$$
$$(\mathbf{w}, \mathbf{z}) \in \mathcal{D}_1, \quad (14)$$

where $\mathcal{D}_1$ represents constraints that exclusively involve the first-stage decision variables w and z.

The solution of the Lagrangian dual provides a lower bound for the model. We enforce the minimum up and down times on slow units in the first-stage problem $(P1)$ in order to recover feasibility at each iteration of the Lagrangian relaxation algorithm. Given these unit commitment schedules, we can solve a second-stage economic dispatch model (ED_s), which is the unit commitment model for scenario s, (UC_s), with u_{gst}, v_{gst} fixed for $g \in G_s$. This choice of decomposition provides feasible solutions at each iteration as well as an upper bound that can be used for computing the duality gap and terminating the algorithm. The algorithm is parallelized both in the solution of (UC_s), as well as the solution of (ED_s), as indicated in Fig. 1. By introducing load shedding in the model as a generating resource with a fuel cost equal to the value of lost load, we ensure that all subproblems (UC_s) and
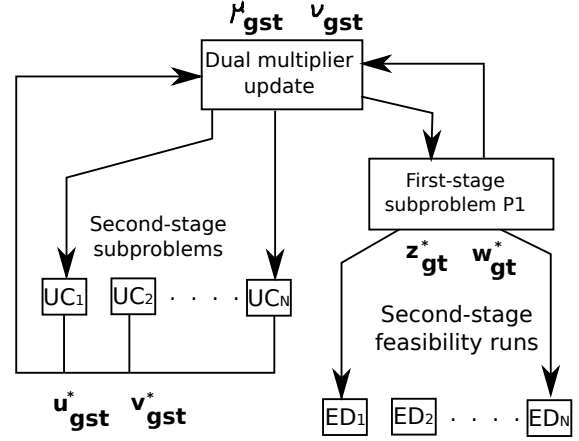

Fig. 1. The parallel implementation of the Lagrangian relaxation algorithm.

(ED_s) are feasible. All sub-problems, $(UC_s), (ED_s), (P1)$, are solved using the branch-and-bound mixed integer linear programming solver of CPLEX. The algorithm is described in further detail in Papavasiliou et al. [15] and Papavasiliou and Oren [14].

In the parallel implementation on the High Performance Computing facility at the Lawrence Livermore National Laboratory we use the subgradient algorithm for updating dual multipliers with Fisher's dual multiplier updating rule [27]. The algorithm is terminated when the desired duality gap has been attained. We have also experimented with a cutting plane algorithm for approximating the dual function, however the performance of the cutting plane algorithm proved to be poor for our problem. In particular, the dual multipliers exhibited oscillations between the box constraints across iterations. Poor performance of the cutting plane algorithm for stochastic unit commitment has also been reported by Takriti et al. [28]. In Section V-B we compare the performance of the subgradient method to bundle algorithms and progressive hedging for the proposed relaxation as well as the relaxation used by Birge [29], [30].

IV. Scenario Selection

The scenario selection algorithm that we describe in this section is inspired by importance sampling and described in detail in Papavasiliou and Oren [14]. Importance sampling is a statistical technique for reducing the number of Monte Carlo simulations that are required for estimating the expected value of a random variable within a certain accuracy. For an exposition see Mazumdar [31] and Infanger [32]. As Pereira and Balu [33] report, this technique has been used in reliability analysis in power systems with composite generation and transmission line failures.

Given a sample space Ω and a measure $\mathbb{P}$ on this space, importance sampling defines a measure $\mathbb{Q}$ on the space that reduces the variance of the observed samples of the random variable C, and weighs each simulated outcome ω by $\mathbb{P}(\omega)/\mathbb{Q}(\omega)$ in order to un-bias the simulation results. The measure $\mathbb{Q}$ is ideally chosen such that it represents the

contribution of a certain outcome to the expected value that is being computed, i.e.

$$\mathbb{Q}^{\star}(\omega) = \frac{\mathbb{P}(\omega)C(\omega)}{\mathbb{E}_{\mathbb{P}}C} \qquad (15)$$

where $\mathbb{E}_{\mathbb{P}}$ denotes the expectation operator with respect to the measure $\mathbb{P}$. Of course, it is not possible to determine this measure since $\mathbb{E}_{\mathbb{P}}C$ is the quantity we wish to compute. Nevertheless, the intuition of selecting samples according to their contribution to the expected value can be carried over to scenario selection.

The extension of the intuition of importance sampling to the case of scenario selection is straightforward: if the ideal measure $\mathbb{Q}^{\star}$ of Eq. (15) were closely approximated by a measure $\mathbb{Q}$, then selecting a small number of outcomes according to this measure and weighing them according to $\mathbb{P}(\omega)/\mathbb{Q}(\omega)$ would provide an accurate estimate of the expected cost. Therefore, samples selected according to $\mathbb{Q}$ can be interpreted as representative scenarios that need to be weighted according to $\mathbb{P}(\omega)/\mathbb{Q}(\omega)$ relative to each other in the objective function of the stochastic unit commitment model in order not to bias the result.

We proceed by generating an adequately large subset of the sample space $\Omega_S = \{\omega^1, \ldots, \omega^M\}$ and we calculate the cost of each sample against a deterministic unit commitment policy $C_D(\cdot)$. Since $\bar{C} = \sum_{i=1}^{M} \frac{C_D(\omega_i)}{M}$ provides an accurate estimate of expected cost, we interpret the sample space of the system as Ω_S and the original measure p as the uniform distribution over Ω_S, hence $\mathbb{P}(\omega) = M^{-1}$ for all $\omega \in \Omega_S$. We then obtain $\mathbb{Q}(\omega_i) = C_D(\omega_i)/(M\bar{C}), i = 1, \ldots, M$, and each selected scenario is weighed according to $\pi_s = \mathbb{P}(\omega)/\mathbb{Q}(\omega)$, hence $\pi_s/\pi_{s'} = C_D(\omega^{s'})/C_D(\omega^s)$ for each pair of selected scenarios $\omega^s, \omega^{s'} \in \hat{\Omega}$. Hence, the proposed algorithm selects scenarios with a likelihood that is proportional to their cost impact, and discounts these scenarios in the stochastic unit commitment in proportion to their cost impact in order not to bias the stochastic unit commitment policy. We therefore arrive at the following algorithm:

Step (a). Define the size N of the reduced scenario set $\hat{\Omega} = \{\omega^1, \ldots, \omega^N\}$.

Step (b). Generate a sample set $\Omega_S \subset \Omega$, where $M = |\Omega_S|$ is adequately large. Calculate the cost $C_D(\omega)$ of each sample $\omega \in \Omega_S$ against the best deterministic unit commitment policy and the average cost $\bar{C} = \sum_{i=1}^{M} \frac{C_D(\omega_i)}{M}$.

Step (c). Choose N scenarios from Ω_S, where the probability of picking a scenario ω is $C_D(\omega)/(M\bar{C})$.

Step (d). Set

$$\pi_s = \frac{C_D(\omega)^{-1}}{\sum_{\omega \in \hat{\Omega}} C_D(\omega)^{-1}}, \omega^s \in \hat{\Omega} \qquad (16)$$

The scenario selection and probability assignment procedure is shown in Fig. 2. We note from the figure that the com-
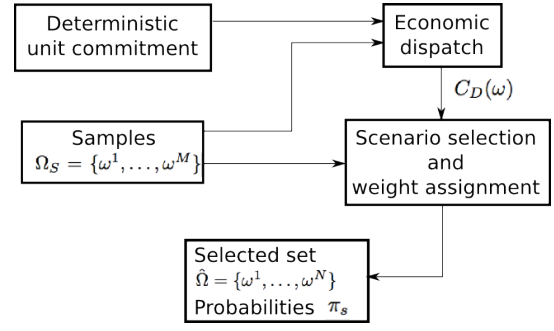

Fig. 2. A schematic of the scenario selection and probability assignment procedure.

putation required for the scenario generation and probability assignment procedure requires the solution of a single deterministic unit commitment problem and a set of M economic dispatch problems. Since the commitment of slow units is fixed in the economic dispatch model, the economic dispatch model solves much faster than the deterministic unit commitment model. The computation time required to solve a collection of economic dispatch problems is negligible, especially since it can be performed in parallel.

An alternative scenario selection algorithm that can be used in two-stage stochastic programming with integer first-stage decision variables is the sample average approximation (SAA) algorithm. The SAA algorithm [18] assigns an equal weight to a finite sample of independently generated scenarios. Intuition suggests that the larger the sample size used in the SAA algorithm, the better the approximation of the original optimization problem and therefore the more likely that the resulting solution to the SAA model will be equal to the true optimum. Indeed, Kleywegt et al. [18] establish that the probability that the optimal solution of the SAA problem will be an optimal solution for the true optimal solution converges to one at an exponential rate in the number of scenarios considered in the SAA problem.

In the results section we present an SAA formulation with 1000 scenarios. The solution of this problem is only possible in a distributed computation environment. We compare the results of the SAA model to the importance sampling algorithm proposed in the previous section.

V. RESULTS

In this section we analyze a test system of the California Independent System Operator interconnected with the Western Electricity Coordinating Council. The system is composed of 225 buses, 375 lines and 130 generators. The horizon of the problem is chosen to be 24 hours, with hourly increments. This is the horizon of the integrated forward market model solved by the California ISO.

We focus on the integration of wind power in this study. The wind penetration level that we analyze corresponds to the 2030 wind integration targets of California. The wind data is calibrated against one year of data from the National Renewable Energy Laboratory database. In order to generate scenarios of wind power production for the simulation we

have developed a time series model of wind power production [34]. Due to the fact that the unit commitment model accounts for transmission constraints, the location of renewable supply needs to be explicitly accounted for in the time series model. The multiregional wind power production model that we use extends the methods of existing literature on wind power time series modeling [35], [36], [37]. The NREL data set provides both wind speed as well as wind power production measurements. Due to the highly non-linear relation between wind speed and wind power, we use an autoregressive model for wind speed and use an aggregate power curve for each location under consideration in order to convert the wind speed time series to a wind power time series. Spatial correlations are accounted for in the wind speed time series model through a covariance matrix in the underlying noise vector that is used for generating the wind speed time series in each location.

Generators and transmission lines are assumed to fail independently with a probability of 1% and 0.1% respectively. The specific failure rates are based on standard values that are assumed in the academic literature [33], [38] as well as generating availability data systems (GADS) for existing power systems [39].

We study four day types corresponding to a weekday for each season. This approach has been adopted by other authors in the literature as well [7] and enables us to simulate a large number of independent Monte Carlo outcomes in the economic dispatch model. This is necessary in order to increase the confidence of our cost results. In contrast, other models that simulate operations sequentially do so for a single year of operations [40], [41]. In the case of the WECC system that we are studying, we have observed that a single year of simulations does not suffice for the moving average cost to converge.

In the following study we present results for the SAA scenario selection algorithm based on 1000 scenarios (denoted as $S1000$), as well as our proposed scenario selection algorithm based on 10, 50 and 100 scenarios (denoted respectively as $S10$, $S50$ and $S100$). This results in four optimization models for each day type. We then solve each of these optimization models for four different duality gaps: 1%, 1.5%, 2% and 2.5%. Intuitively, we expect that an increasing number of scenarios will result in an improved policy, albeit with increased computational effort. On the other hand, a decreasing duality gap is also expected to deliver superior performance, again at the cost of increased computational effort.

The target optimality gap of 1% for the optimization problem dictates a smaller step size for the subgradient method. We have set a step size of $\lambda = 0.5\%$ in Fisher's updating rule [15]. We have set the MIP gap for the solution of (ED_s) to 0.1% and the MIP gap for the solution of $(P1)$ and $(P2_s)$ to 0.5%. Using a larger MIP gap for the solution of $(P1)$ and $(P2_s)$ reduces the solution time of these problems, however this introduces a larger error to the overall optimality gap. In order to mitigate load imbalance, we have enforced a time limit of 60 seconds for the solution of each subproblem $(P2_s)$ which in most instances suffices to achieve the target MIP gap of 0.5% for $(P2_s)$.

The simulations have been performed on the Cab cluster of the Lawrence Livermore National Laboratory. The Cab cluster consists of 1296 nodes with 20736 cores, equipped with an Intel Xeon E5-2670 processor at 2.6GHz and 32 GB per node. We have used CPLEX 12.4 through the CPLEX Java callable library. MPI was used for parallelizing the application.

The stochastic unit commitment model with 1000 scenarios has 3,121,800 binary variables, 20,643,120 continuous variables, and 66,936,000 constraints. By comparison, the WILMAR model of Tuohy et al. [40], which is a state-of-the-art stochastic unit commitment model with a 36-hour horizon, has 179,000 constraints and 167,000 variables of which 16,000 are integer. A deterministic instance of the Pennsylvania Jersey Maryland market with a 24-hour horizon includes 24,264 binary variables, 833,112 continuous variables and 1,930,776 constraints. We therefore note that the size of the problem addressed in this paper is one to two orders of magnitude greater than state-of-the-art stochastic unit commitment models and deterministic models of industrial scale.

A. Cost Performance

An interesting practical tradeoff that we can explore by exploiting high performance computing is the relative influence of scenario size versus optimality gaps on the cost performance of the resulting policy. A larger number of scenarios results in a more accurate representation of uncertainty in the model, however it requires more computation at each iteration of the Lagrangian relaxation algorithm. Therefore, the question rises whether it is worth sacrificing the accuracy with which we represent uncertainty in our model by using a smaller set of scenarios, in order to achieve an improved duality gap.

In order to address this question, we present a ranking of the costs for each policy in Fig. 3. The policies are ranked in order of increasing cost. The error bars indicate a 69.5% confidence in the estimation of the average cost, namely they are computed as the interval $(\bar{C} - \frac{\hat{S}_N}{\sqrt{N}}, \bar{C} + \frac{\hat{S}_N}{\sqrt{N}})$, where

$$\bar{C} = \frac{1}{N} \sum_{n=1}^{N} C_n \tag{17}$$

$$\hat{S}_n^2 = \frac{\sum_{n=1}^{N} (C_n - \bar{C})^2}{N - 1}. \tag{18}$$

The results of Fig. 3 suggest that the reduction of the duality gap can be as influential towards the performance of the algorithm as the increased number of scenarios. Policies derived from the smallest duality gaps appear as the top performers, whereas even policies with 1000 scenarios may perform badly if the duality gap is relatively high. For example, among the three worse policies in summer we find the models with 1000 scenarios for a gap of 2% and 2.5%. The best policy for all day types has a 1% optimality gap (but only the largest number of scenarios for spring), and for all but one day type the worst policy has a 2.5% gap. For spring, the best performing policy is the one with the greatest number of scenarios and smallest gap ($G = 1, S = 1000$). For spring, summer and fall the worst policy is the one with the fewest scenarios and the greatest gap, namely $G = 2.5, S = 10$.
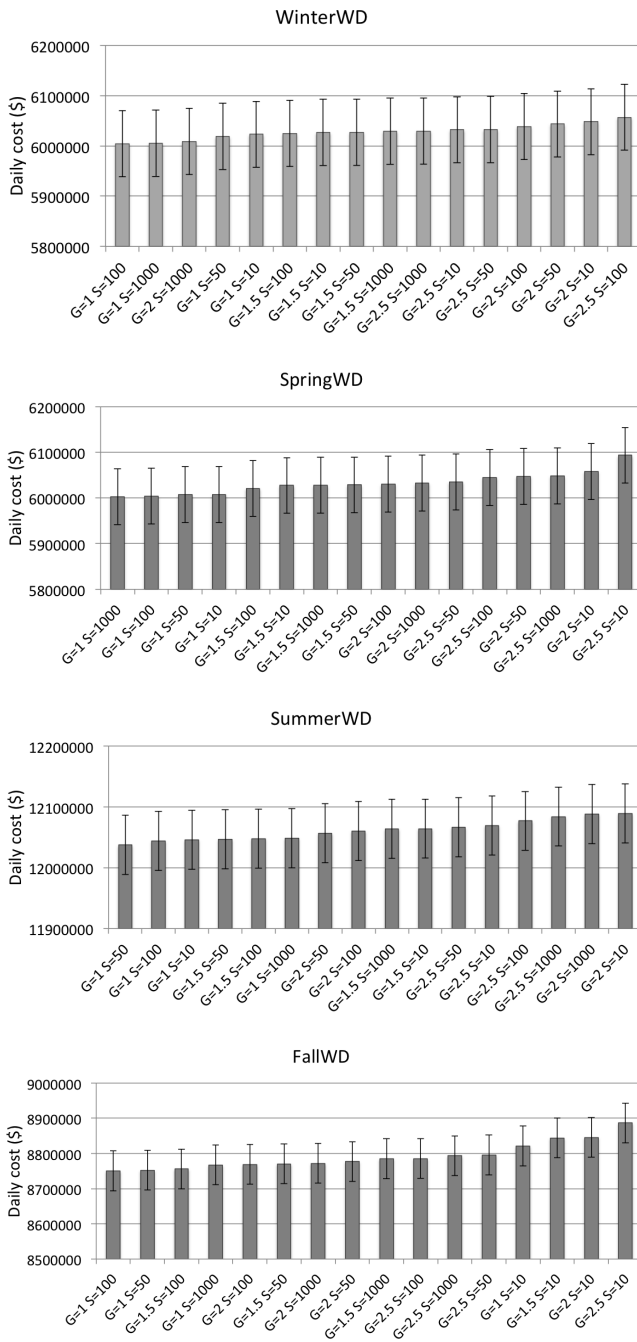
Fig. 3. The cost of each policy for each day type.

Fig. 4. The daily unit commitment schedule for the best, average and worse unit commitment policy for each day type.

The confidence intervals are weakest for winter, in the sense that the worst policy is within the confidence interval of the best one. Instead, for fall and spring the worst policy is well outside the confidence interval of the best policy. The fact that the differences among policies often fall within the confidence intervals of the simulation supports the approach that we have taken in using day types, which permits parallel and independent simulation of days, thereby increasing the sample size.

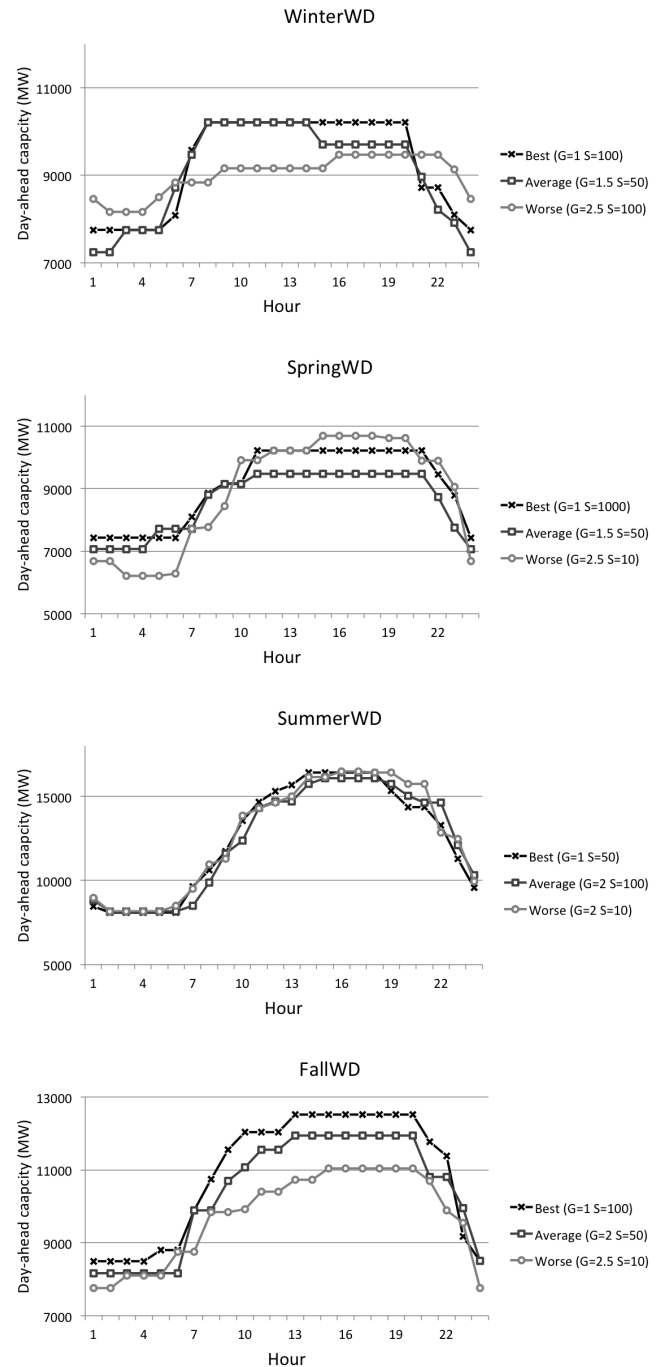The results of Fig. 3 also validate the satisfactory performance of the proposed scenario selection algorithm. Note that the top performance for winter, summer and fall is attained by a policy derived from the proposed scenario selection algorithm based on importance sampling, rather than the SAA algorithm, despite the fact that the SAA algorithm utilizes at least ten times more scenarios. For all day types, the importance sampling algorithm results in a policy that is within the top 2 performers. An important practical implication of these results is that satisfactory performance can be attained by models of moderate scale, provided an appropriate scenario
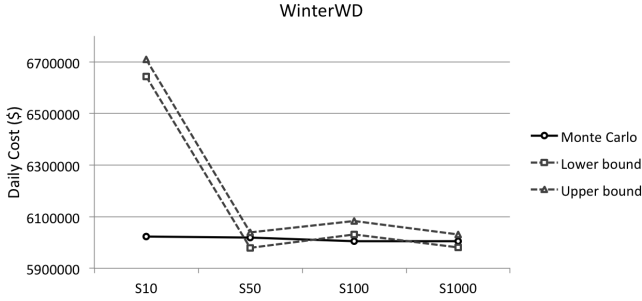
Fig. 5. Lower bound, upper bound and Monte Carlo performance for each scenario for winter weekdays (1% optimality gap).

selection policy is utilized. High performance computing may not be necessary for solving moderately sized stochastic unit commitment models within an operationally acceptable time frame, which could accelerate the adoption of stochastic programming in short-term scheduling. Nevertheless our ability to arrive at this conclusion by comparing reduced scenario policies with a 1000-scenario policy was only made possible by solving a problem whose scale could not have been addressed through serial computational methods.

The amount of day-ahead reserve capacity that is committed is shown in Fig. 4. In each panel we present the best policy, average policy (ranked 8th in terms of cost) and worst policy. For winter we observe that the worst policy is significantly under-committing resources in the middle of the day and over-committing in the beginning and end of the day. For spring the worst unit commitment varies relative to the best policy throughout the day, under-committing in the beginning of the day and over-committing during peak load. In summer the worst policy tends to over-commit, especially around 7 pm. In fall the worse policy tends to under-commit.

The Monte Carlo cost and upper bound are evaluated for the same first-stage decision against the sample set and scenario set respectively. Therefore we expect that these two quantities will converge as the number of scenarios increases. In Fig. 5 we present the lower bound and upper bound of the Lagrangian relaxation algorithm as well as the estimated Monte Carlo cost as a function of the number of scenarios for instances with a target gap of 1% for winter weekdays. We indeed observe in Fig. 5 that the upper bound converges to the Monte Carlo cost. The corresponding graphs for the other day types are omitted in order not to overburden the presentation. The results for summer follow those of winter. For spring, the upper bound estimate is most accurate for 50 scenarios. For fall, the initial deviation between upper bound and Monte Carlo estimate is quite large, and is minimized at 100 scenarios.

The strong influence of reducing the duality gap on the performance of the stochastic unit commitment policy is reinforced by results obtained by testing an alternative scenario selection algorithm, proposed by Morales et al. [42], that accounts for the impact of scenarios on the objective function. The algorithm proposed in [42] was implemented in this study by solving a deterministic equivalent unit commitment problem in which hourly renewable power production is set to

its average over all candidate scenarios in Eq. (2), the capacity of all generators is de-rated by their average availability (0.99) in Eq. (3), and the capacity of all line susceptances is de-rated by their average availability (0.999) in Eq. (4). This unit commitment policy is used for quantifying the impact of each scenario on the objective function, which enables us to select scenarios and assign probabilities to the selected scenarios. The scenario selection algorithm proposed by Morales et al. [42] was implemented for 10 scenarios and a gap of 1% ($G = 1, S = 10$). The resulting unit commitment policy ranked ninth, fifth, first and fourth for Winter, Spring Summer and Fall respectively, relative to the unit commitment policies studied in this paper and shown in Fig. 3. For two instances (Summer and Fall), the policy outperforms the policy generated by the importance sampling algorithm with $G = 1, S = 10$.

B. Running Times

The running times of the policies are shown in Fig. 6. The reported run times correspond to using a number of cores equal to the number of scenarios for each problem instance, and therefore correspond to the best possible run times one can attain. We note that all policies are computed within 24 hours, given a sufficient number of cores is available. This implies that the proposed approach is implementable within the day-ahead operating time frame. Moreover, we observe that certain policies that are easy to compute can perform quite well. For example, the 10-scenario policy with a 1% optimality gap is the third best policy for summer weekdays. Run times are most regular for fall and summer, in the sense that the most time-consuming policies are the best-performing ones, and vice versa.

In Fig. 7 we plot the dependency of running times on the number of processor cores used for winter weekdays. Results for the other day types lead to similar conclusions and have been omitted in order to facilitate the communication of our conclusions. Each problem instance has been solved on 10, 50, 100 and 1000 cores. As we describe in Fig. 1, the decomposition is parallelized across scenarios. Therefore, we did not solve a certain instance with more cores than the number of scenarios in the model. For example, problems with 50 scenarios were only run for 10 and 50 cores, not 100 or 1000 cores. All figures range between 0 and 48 hours of computation. Some of the measurements are well outside the range of these charts. Certain problem instances were interrupted after 5 days of run time without terminating within the target duality gap, however given a sufficient number of processors all problem instances were solved within 24 hours as we demonstrate in Fig. 6.

The results suggest that the speedup achieved through the parallelization of computation can reduce computation times to an operationally acceptable time frame. To illustrate this fact, note that the solution of the winter $G = 1, S = 1000$ instance can be sped up from more than 120 hours (5 days) with 10 processors to 17.5 hours when 1000 processors are available. The influence of the target duality gap on running time is also evident in these figures and represents an important tradeoff between the quality of the resulting solution and the available computation time.
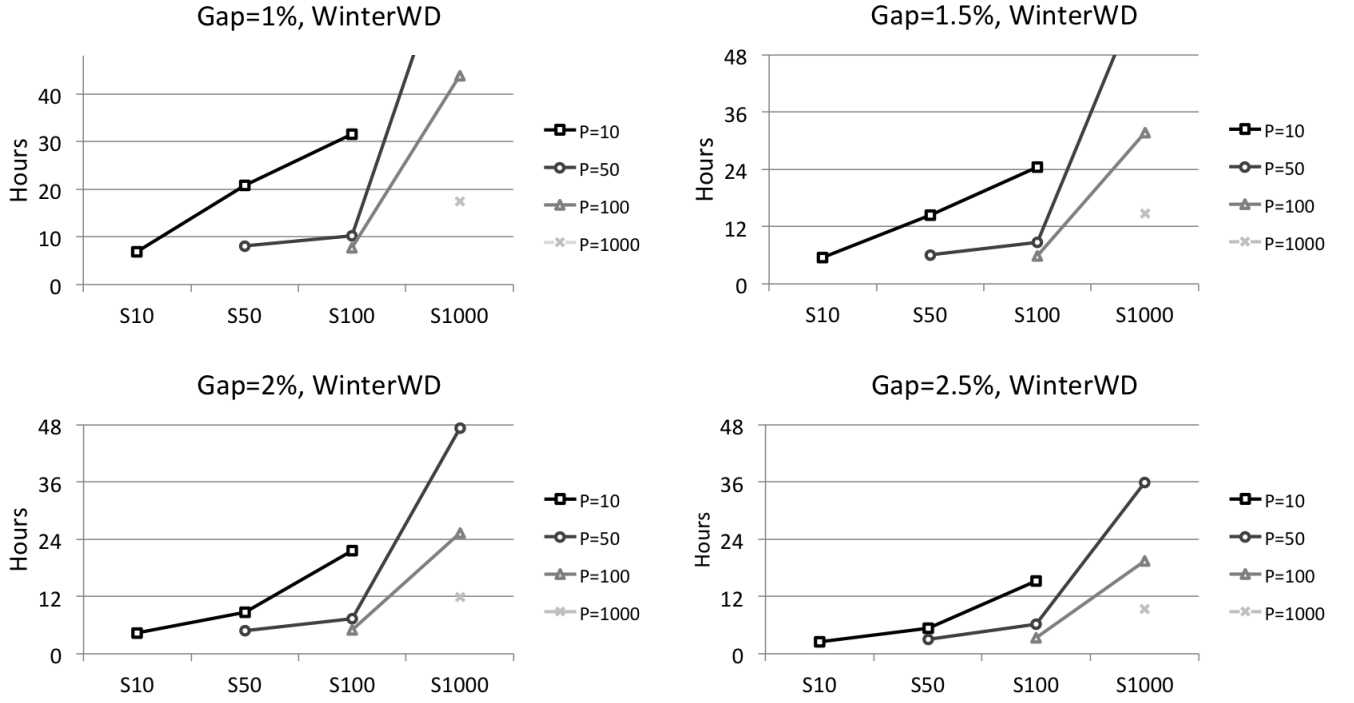
Fig. 7. The run times for varying number of processor cores for Winter weekdays.

Figs. 3, 6 and 7 indicate that the number of scenarios that should be used in operating practice for short-term stochastic scheduling depends on the amount of available computation time and the number of available computational resources. Regardless of the number of scenarios that are input to the problem, it is commonly preferable to exhaust the available computing time in order to decrease the duality gap as much as possible. A smaller duality gap corresponds to a smaller violation of the non-anticipativity constraints, or equivalently a solution that is closer to the set of day-ahead feasible schedules. Although there is no theoretical guarantee that a smaller gap for the same instance will deliver a better result (compare, for example, the case of $G = 2$ with the case of $G = 2.5$ for $S = 10$ for winter weekdays in Fig. 3), the computational experiments that are reported here suggest that this tends to be the case.

A further indication about the promising contribution of parallel computing in day-ahead power system operations under uncertainty is the fact that the running times presented in Fig. 7 could be outperformed by alternative relaxations and dual function optimization methods. In Fig. 8 we present the lower bound derived for the optimization problem (SUC) derived within the 50 first iterations for five alternatives:

- (SG1): The Lagrange relaxation presented in Section III, with the subgradient algorithm and a step-size of $\lambda = 1\%$ in Fisher's multiplier updating rule [27]. This is the reference case against which we compare the following alternative implementations.
- (SG2): Lagrange relaxation of the unit commitment non-anticipativity constraints using a single equality constraint, as in [29], [30]:

$$(1 - \pi_{s_0})u_{gs_0t} - \sum_{s \in S - \{s_0\}} \pi_s u_{gst} = 0, (\mu_{gt}) \quad (19)$$

The subgradient algorithm is used for dual function optimization, with $\lambda = 0.1\%$ in Fisher's multiplier updating rule. The advantage of this approach is the fact that the dual space is of smaller dimension by a factor of $|S|$ (the cardinality of the scenario set). However, this relaxation may result in inferior bounds due to the smaller number of dual variables that can be used for coordinating per-scenario decisions and is not guaranteed to provide a feasible solution, and therefore an upper bound, at each iteration.

- (B1): A bundle algorithm [43] applied to the Lagrangian relaxation of Section III. Ascent steps are taken whenever the dual function, evaluated at the new iterate, is better than the dual function evaluated at the incumbent center of the quadratic term. The coefficient of the quadratic term is set so that whenever an ascent step takes place the coefficient is decreased by a factor of 5, whereas when a null step takes place the coefficient is increased by a factor of $0.2^{0.2}$, i.e. it takes 5 null steps to "lose" the confidence that we gained in our cutting plane model from one ascent step.
- (B2): A bundle algorithm applied to the relaxation used by Birge [29], [30]. As in the case of (SG2), the algorithm does not necessarily generate upper bounds at any given iteration. The coefficient of the quadratic term is set so that the whenever an ascent step takes place the coefficient is decreased by a factor of 5, whereas when
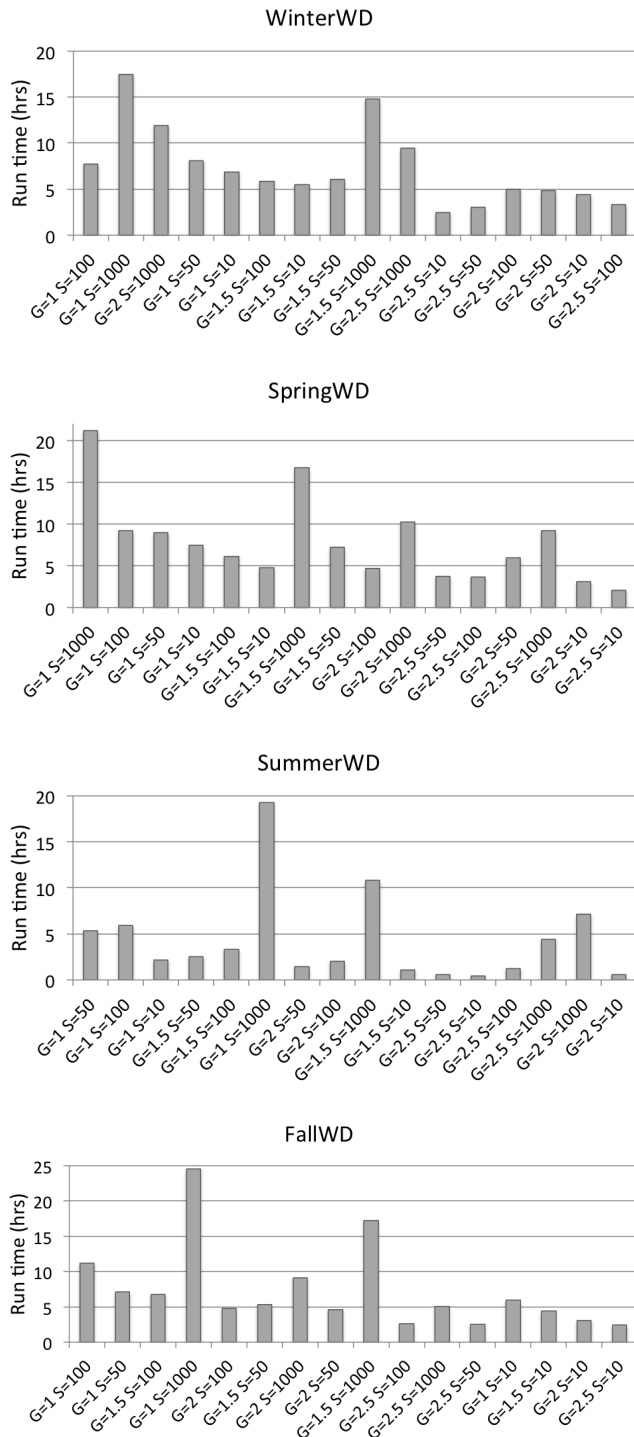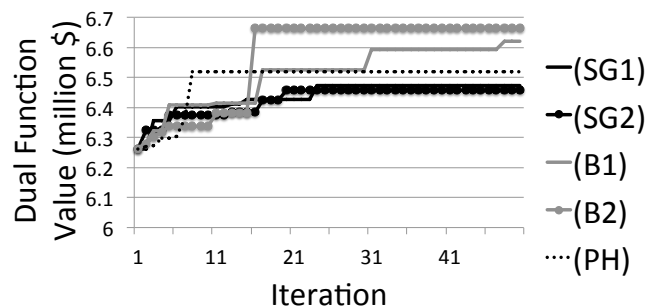
Fig. 8. The lower bounds of (SUC) for 50 iterations of $(SG1)$, $(SG2)$, $(B1)$, $(B2)$ and (PH) for winter weekdays.

a null step takes place the coefficient is increased by a factor of $0.2^{0.1}$, i.e. it takes 10 null steps to "lose" the confidence that we gained in our cutting plane model from one ascent step.

- (PH): A progressive hedging algorithm [2], [6], [44], [45]. The progressive hedging algorithm is not guaranteed to yield a lower bound or feasible solution in any of its iterations, therefore we additionally employ the lower bound proposed in [46]. We use a common coefficient ρ for the quadratic term, which is a fraction (10%) of the average minimum load cost, in accord with the prescribed heuristic for setting ρ by Watson and Woodruff [45]. We do not consider a progressive hedging algorithm applied to the relaxation used by Birge [29], [30], since it merely results in a rescaling of the relaxed non-anticipativity constraint.

The important common feature of all the proposed alternatives to $(SG1)$ is that they too can be parallelized by scenario. Moreover, all alternatives with the exception of $(SG2)$ appear to achieve faster convergence, while the performance of $(SG2)$ is comparable to that of $(SG1)$. This indicates that the running times of Fig. 7 can be improved upon. This reinforces the potential benefits of parallel computing in short-term power system operations. The detailed comparison and analysis of alternative relaxations and algorithms in a parallel computing environment presents a promising area of future research.

VI. CONCLUSIONS AND PERSPECTIVES

In this paper we have presented an application of high performance computing for the study of the stochastic unit commitment problem. We believe that high performance computing presents an exciting opportunity for the integration of stochastic programming in an operational setting in order to address hourly renewable supply uncertainty in short-term power systems scheduling, analogously to the proliferation of stochastic programming in medium-term hydrothermal scheduling in order to address the monthly uncertainty of precipitation [41].

The results of this paper validate the performance of the stochastic unit commitment policy, as we are able to attain a performance which is comparable to the SAA algorithm using a number of scenarios which is 10-100 smaller. We have also demonstrated that high performance computing can

Fig. 6. The run time of each policy for each day type.

reduce the computation times of stochastic unit commitment to a level which is acceptable from an operational perspective. A comparison of the subgradient algorithm to bundle methods and progressive hedging reinforces this conclulsion since these methods can also be decomposed by scenario and exhibit faster convergence than the subgradient method, as theoretical analysis and computational experience suggest. Finally, we have explored the relative importance of scenario set size and duality gaps on performance. We conclude that reducing the duality gap of the Lagrangian relaxation yields comparable benefits to increasing the size of the scenario set. This indicates, form a practitioner perspective, that stochastic unit commitment can become viable with moderately sized instances provided a reliable scenario selection algorithm is available.

The present research has raised numerous interesting areas of future research. We are interested in further analyzing alternative relaxations of the non-anticipativity constraints [29] and comparing the relative performance of alternative distributed dual function optimization methods. In future research we also intend to explore issues relating to threading, and other aspects of performance tuning. An area of theoretical research that warrants further investigation is what feature of the stochastic unit commitment problem causes a non-zero duality gap. Our ultimate goal is to utilize high performance computing in order to solve industrial scale stochastic unit commitment problems. Current efforts are focused on solving the Pennsylvania Jersey Maryland model that consists of 13,867 buses, 1,011 generators and 18,824 branches, and is made available by the Federal Energy Regulatory Commission. In solving large unit commitment problems in practice (e.g. PJM) it is common practice to focus on a reduced subset of constraints that are monitored and ignore the others (which is possible when the constraints are characterized in terms of a PTDF formulation). This significantly reduces the effective problem size and this is an approach that we will adopt in our future work. Promising future applications of high performance computing include topology control in short-term scheduling operations, alternative approaches to reserve scheduling such as robust optimization, as well as investment planning models.

APPENDIX

A. Notation

Sets

G: set of all generators, G_s: subset of slow generators, G_f: subset of fast generators

G_n: set of generators that are located in bus n

S: set of scenarios, T: set of time periods, L: set of lines, N: set of nodes

Decision variables

u_{gst}: commitment, v_{gst}: startup, p_{gst}: production of generator g in scenario s, period t

θ_{nst}: phase angle at bus n in scenario s, period t

w_{gt}: commitment, z_{gt}: startup of slow generator g in period t

e_{lst}: power flow on line l in scenario s, period t

Parameters

π_s: probability of scenario s

K_g: minimum load cost, S_g: startup cost, C_g: marginal cost of generator g

D_{nst}: demand in bus n, scenario s, period t

P_{gs}^+, P_{gs}^-: minimum and maximum capacity of generator g in scenario s

B_{ls}: susceptance of line l in scenario s

TC_l: capacity of line l

REFERENCES

[1] R. Piwko, K. Clark, L. Freeman, G. Jordan, and N. Miller, "Western wind and solar integration study," National Renewable Energy Laboratory, Tech. Rep., May 2010.

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

[3] P. Carpentier, G. Cohen, J.-C. Culioli, and A. Renaud, "Stochastic optimization of unit commitment: a new decomposition framework," *IEEE Transactions on Power Systems*, vol. 11, no. 2, pp. 1067–1073, May 1996.

[4] J. A. Muckstadt and S. A. Koenig, "An application of lagrangian relaxation to scheduling in power-generation systems," *Operations Research*, vol. 25, no. 3, pp. 387–403, May-June 1977.

[5] M. P. Nowak and W. Römisch, "Stochastic lagrangian relaxation applied to power scheduling in a hydro-thermal system under uncertainty," *Annals of Operations Research*, vol. 100, no. 1-4, pp. 251–272, 2000.

[6] S. Takriti and J. R. Birge, "Lagrangian solution techniques and bounds for loosely coupled mixed-integer stochastic programs," *Operations Research*, vol. 48, no. 1, pp. 91–98, January - February 2000.

[7] D. Bertsimas, E. Litvinov, X. A. Sun, J. Zhao, and T. Zheng, "Adaptive robust optimization for the security constrained unit commitment problem," *IEEE Transactions on Power Systems*, vol. 28, no. 1, pp. 52–63, February 2013.

[8] L. Wu, M. Shahidehpour, and T. Li, "Stochastic security-constrained unit commitment," *IEEE Transactions on Power Systems*, vol. 22, no. 2, pp. 800–811, May 2007.

[9] J. Wang, M. Shahidehpour, and Z. Li, "Security-constrained unit commitment with volatile wind power generation," *IEEE Transactions on Power Systems*, vol. 23, no. 3, pp. 1319–1327, August 2008.

[10] U. A. Ozturk, M. Mazumdar, and B. A. Norman, "A solution to the stochastic unit commitment problem using chance constrained programming," *IEEE Transactions on Power Systems*, vol. 19, no. 3, pp. 1589–1598, August 2004.

[11] J. Dupacova, N. Gröwe-Kuska, and W. Römisch, "Scenario reduction in stochastic programming: An approach using probability metrics," *Math Programming*, vol. 95, no. 3, pp. 493–511, 2003.

[12] H. Heitsch and W. Römisch, "Scenario reduction algorithms in stochastic programming," *Computational Optimization and Applications*, vol. 24, no. 2-3, pp. 187–206, 2003.

[13] N. Gröwe-Kuska, K. C. Kiwiel, M. P. Nowak, W. Romisch, and I. Wegner, *Decision Making Under Uncertainty: Energy and Power*, ser. IMA Volumes in Mathematics and Its Applications. New York: Springer-Verlag, 2002, vol. 128, ch. Power Management in a Hydro-Thermal System Under Uncertainty by Lagrangian Relaxation, pp. 39–70.

[14] A. Papavasiliou and S. S. Oren, "Multi-area stochastic unit commitment for high wind penetration in a transmission constrained network," *Operations Research*, vol. 61, no. 3, pp. 578–592, May / June 2013.

[15] A. Papavasiliou, S. S. Oren, and R. P. O'Neill, "Reserve requirements for wind power integration: A scenario-based stochastic programming framework," *IEEE Transactions on Power Systems*, vol. 26, no. 4, pp. 2197–2206, November 2011.

[16] K. Hoyland, M. Kaut, and S. W. Wallace, "A heuristic for moment-matching scenario generation," *Computational Optimization and Applications*, vol. 24, pp. 169–185, 2003.

[17] J. M. Latorre, S. Cerisola, and A. Ramos, "Clustering algorithms for scenario tree generation: Application to natural hydro inflows," *European Journal of Operational Research*, vol. 181, pp. 1339–1353, 2007.

[18] A. J. Kleywegt, A. Shaprio, and T. Homem-De-Mello, "The sample average approximation method for stochastic discrete optimization," *SIAM Journal on Optimization*, vol. 12, no. 2, pp. 479–502, 2001.

[19] A. Monticelli, M. V. F. Pereira, and S. Granville, "Security-constrained optimal power flow with post-contingency corrective rescheduling," *IEEE Transactions on Power Systems*, vol. 2, no. 1, pp. 175–180, February 1987.

[20] M. V. F. Pereira, M. J. Teixeira, M. F. McCoy, and H. J. C. P. Pinto, "Developing concurrent processing applications to power system planning and operations," *IEEE Transactions on Power Systems*, vol. 5, no. 2, pp. 659–664, May 1990.

[21] B. H. Kim and R. Baldick, "Coarse-grained distributed optimal power flow," *IEEE Transactions on Power Systems*, vol. 12, no. 2, pp. 932–939, May 1997.

[22] A. G. Bakirtzis and P. N. Biskas, "A decentralized solution to the DC-OPF of interconnected power systems," *IEEE Transactions on Power Systems*, vol. 18, no. 3, pp. 1007–1013, August 2003.

[23] P. N. Biskas, A. G. Bakirtzis, N. I. Macheras, and N. K. Pasialis, "A decentralized implementation of DC optimal power flow on a network of computers," *IEEE Transactions on Power Systems*, vol. 20, no. 1, pp. 25–33, February 2005.

[24] D. Falcao, *High Performance Computing in Power System Applications*, ser. Vector and Parallel Processing - VECPAR'96, Second International Conference, Selected Papers, Porto, Portugal, September 25-27, J. M. L. M. Palma and J. Dongarra, Eds. Springer, 1997, vol. 1215.

[25] E. M. Constantinescu, V. M. Zavala, M. Rocklin, S. Lee, and M. Anitescu, "A computational framework for uncertainty quantification and stochastic optimization in unit commitment with wind power generation," *IEEE Transactions on Power Systems*, vol. 26, no. 1, pp. 431–441, February 2011.

[26] P. A. Ruiz, R. C. Philbrick, and P. W. Sauer, "Wind power day-ahead uncertainty management through stochastic unit commitment policies," in *Power Systems Conference and Exposition*, March 2009, pp. 1–9.

[27] M. L. Fisher, "An applications oriented guide to Lagrangian relaxation," *Interfaces*, vol. 15, no. 2, pp. 10–21, 1985.

[28] S. Takriti, B. Krasenbrink, and L. S.-Y. Wu, "Incorporating fuel constraints and electricity spot prices into the stochastic unit commitment problem," *Operations Research*, vol. 48, no. 2, pp. 268–280, March - April 2000.

[29] J. R. Birge and M. A. H. Dempster, "Stochastic programming approaches to stochastic scheduling," *Journal of Global Optimization*, vol. 9, no. 3-4, pp. 417–451, 1996.

[30] S. Sen, J. L. Higle, and J. R. Birge, "Duality gaps in stochastic integer programming," *Journal of Global Optimization*, vol. 18, pp. 189–194, 2000.

[31] M. Mazumdar, *Importance Sampling in Reliability Estimation, Reliability and Fault Tree Analysis*. Philadelphia, PA: SIAM, 1975, pp. 153–163.

[32] G. Infanger, "Monte carlo (importance) sampling within a Benders decomposition algorithm for stochastic linear programs," *Annals of Operations Research*, vol. 39, no. 1, pp. 69–95, 1992.

[33] M. V. F. Pereira and N. J. Balu, "Composite generation/transmission reliability evaluation," *Proceedings of the IEEE*, vol. 80, no. 4, pp. 470–491, April 1992.

[34] A. Papavasiliou, "Coupling renewable energy supply with deferrable demand," Ph.D. dissertation, U.C. Berkeley, October 2011.

[35] B. G. Brown, R. W. Katz, and A. H. Murhpy, "Time series models to simulate and forecast wind speed and wind power," *Journal of Climate and Applied Meteorology*, vol. 23, pp. 1184–1195, 1984.

[36] J. L. Torres, A. Garcia, M. D. Blas, and A. D. Francisco, "Forecast of hourly wind speed with ARMA models in Navarre (Spain)," *Solar Energy*, vol. 79, no. 1, pp. 65–77, July 2005.

[37] J. M. Morales, R. Minguez, and A. J. Conejo, "A methodology to generate statistically dependent wind speed scenarios," *Applied Energy*, vol. 87, pp. 843–855, 2010.

[38] C. Grigg, P.Wong, P. Albrecht, R. Allan, M. Bhavaraju, R. Billinton, O. Chen, C. Fong, S. Haddad, S. Kuruganty, W. U. R. Mukerji, D. Patton, N. Rau, D. Reppen, K. Schneider, M. Shahidehpour, and C. Singh, "The IEEE reliability test system - 1996," *IEEE Transactions on Power Systems*, vol. 14, no. 3, pp. 1010 – 1020, August 1999.

[39] CEA. (2006, December) Forced outage rates. [Online]. Available: http://www.science.smith.edu/~jcardell/Courses/EGR325/Readings/FOR%202006Update.pdf

[40] A. Tuohy, P. Meibom, E. Denny, and M. O'Malley, "Unit commitment for systems with high wind penetration," *IEEE Transactions on Power Systems*, vol. 24, no. 2, pp. 592–601, May 2009.

[41] M. V. F. Pereira and L. M. V. G. Pinto, "Multi-stage stochastic optimization applied to energy planning," *Mathematical Programming*, vol. 52, pp. 359–375, 1991.

[42] J. M. Morales, S. Pineda, A. J. Conejo, and M. Carrion, "Scenario reduction for futures trading in electricity markets," *IEEE Transactions on Power Engineering*, vol. 24, no. 2, pp. 878–888, May 2009.

[43] C. Lemaréchal, "Lagrangian relaxation," in *Computational Combinatorial Optimization, Lecture Notes in Computer Science*. Springer, 2001, vol. 2241, pp. 112–156.

[44] 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.

[45] J.-P. Watson and D. L. Woodruff, "Progressive hedging innovations for a class of stochastic mixed-integer resource allocation problems," *Computational Management Science*, vol. 8, pp. 355–370, 2011.

[46] D. Gade, S. M. Ryan, J.-P. Watson, and D. L. Woodruff. (2013) A new lower bound from progressive hedging. [Online]. Available: http://dinamico2.unibg.it/icsp2013/doc/ms/smricsp2013.pdf

Anthony Papavasiliou is an assistant professor in the Department of Mathematical Engineering at the Catholic University of Louvain, and a member of the Center for Operations Research and Econometrics. His research interests are focused on energy systems operations, planning and economics, and optimization under uncertainty. He has served as a researcher for Pacific Gas and Electric and the Federal Energy Regulatory Commission and the Palo Alto Research Center.

Shmuel S. Oren received the B.Sc. and M.Sc. degrees in mechanical engineering and in materials engineering from the Technion Haifa, Israel, and the MS. and Ph.D. degrees in engineering economic systems from Stanford University, Stanford, CA, in 1972. He is a Professor of IEOR at the University of California at Berkeley and the Berkeley site director of the Power System Engineering Research Center (PSERC). He has published numerous articles on aspects of electricity market design and has been a consultant to various private and government organizations.

Barry Rountree received a B.A. in Theater Arts and Drama from the Honors Tutorial College at Ohio University, a M.S. in System and Network Administration from Florida State University, and a Ph.D. in Computer Science from the University of Arizona. He is a staff scientist at the Center for Applied Scientific Computing at Lawrence Livermore National Laboratory.