2013 European Control Conference (ECC)
July 17-19, 2013, Zürich, Switzerland.

A Comparative Study of Stochastic Unit Commitment and Security-Constrained Unit Commitment Using High Performance Computing

Anthony Papavasiliou and Shmuel S. Oren

Abstract— **The large-scale integration of renewable resources has recently raised interest in systematic methods for committing locational reserves in order to secure the system against contingencies and the unpredictable and highly variable fluctuation of renewable energy supply, while accounting for power flow constraints imposed by the transmission network. In this paper we compare two approaches for committing locational reserves: stochastic unit commitment and a hybrid approach of scenario-based security-constrained commitment. Parallel algorithms are developed for solving the resulting models, based on Lagrangian relaxation and Benders decomposition. The proposed algorithms are implemented in a high performance computing environment and the performance of the resulting policies is tested against a reduced model of the California ISO interconnected with the Western Electricity Coordinating Council.**

I. Introduction

The increasing uncertainty of power system operations due to the large-scale integration of renewable energy resources and demand response has raised an interest in systematic methods for committing day-ahead reserves in order to operate the system reliably. Traditional reserve commitment approaches [1] rely on reserve requirements and security constraints that are meant to mitigate continuous fluctuations in demand and renewable supply as well as discrete disturbances such as generator and transmission line failures. However, these models often fail to capture the full range of complexity in an uncertain environment and rely instead on heuristic practices adopted by operators through experience.

The power system operations literature has proposed four fundamental paradigms for representing uncertainty and optimizing the commitment of reserves at least cost: stochastic optimization, security-constrained approaches, robust optimization and probabilistic constraints. It is envisioned that these approaches can be used in the reliability unit commitment phase by the system operator in order to make improved decisions on the commitment of reserves. Stochastic programming was originally posed in the context of unit commitment by Takriti and Birge [2] as an approach for mitigating demand uncertainty and generator outages. Subsequently, numerous variants of the stochastic unit commitment model have been proposed [3], [4], [5], [6], [7], [8], [9] that vary based on the number of stages, the source of uncertainty, the representation of uncertainty and solution methods that are used. The drawback of stochastic unit commitment is the requirement to represent uncertainty in a detailed fashion, by using a large number of appropriately weighted scenarios. The resulting problems are large-scale and Lagrangian relaxation is often employed in order to decompose the problem to tractable subproblems.

Security-constrained unit commitment models require that the system be capable of withstanding major element failures without shedding load. Security constraints address discrete failures of network elements, whereas continuous sources of uncertainty are addressed either through scenarios or exogenous reserve criteria. Wang et al. [10] account for supply and demand fluctuations through exogenous reserve criteria [10] and use Bender's decomposition to solve the problem, while Wu et al. [11] account for continuous disturbances through scenarios and use Lagrangian relaxation in order to solve the problem.

Two additional systematic approaches to short-term scheduling under uncertainty that are not explored in this paper include robust optimization and probabilistic constraints. In robust optimization models [12], [13] the objective is to commit reserves in order to minimize the cost of operating the system against the worst-case realization of uncertainty. The worst-case nature of these approaches reflects more closely the tendency of operators to operate the system in a conservative fashion. Probabilistically constrained models [14], [15] commit units in order to ensure the satisfaction of demand within an exogenously defined probability. Both robust optimization and probabilistically constrained models share the advantage that they require only limited information about the process that drives uncertainty in operations.

The methods discussed above increase the computational challenges associated with unit commitment problems. The solution methods used for solving short-term scheduling problems can often be parallelized. As a result, these applications invite the use of distributed computing for addressing the challenge of short-term scheduling under uncertainty.

Distributed computation has a rich history in the area of power systems planning and operations. A review of the application of high performance computing in power systems is presented by Falcao [16]. Pereira et al. [17] present the application of distributed computing in reliability evaluation for composite outages, scenario analysis for hydro dominated systems and security-constrained dispatch. Monticelli et al. [18] formulate the security-constrained optimal power flow with corrective rescheduling and demonstrate economic benefits in the dispatch on an IEEE system with 118 buses. Kim and Baldick [19] present a parallel algorithm for solving distributed optimal power flow. Bakirtzis and Biskas [20] propose a decentralized Lagrangian relaxation algorithm for solving the optimal power flow problem presented by Kim and Baldick [19]. A parallel implementation of the algorithm in Bakirtzis and Biskas [20] using PVM is presented by

Biskas et al. [21].

Despite the fact that there is a rich body of literature focused on short-term scheduling under uncertainty, the relative performance of the models proposed in the literature is not compared adequately in order to appreciate the tradeoffs involved in using each paradigm. This paper serves two purposes. The first objective is to motivate this discussion by presenting a comparative study of stochastic programming and security-constrained unit commitment models. The second objective is to demonstrate the benefits of distributed computation in accelerating the solution of these models. We present our models in Section II and decomposition methods for solving these models in Section III. A case study of the California ISO interconnected with the Western Electricity Coordinating Council is presented in Section IV. We summarize our conclusions in Section V.

II. Model Description

The problem that this paper is focusing on 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 model that we present in this paper accounts for transmission constraints over the set of lines K, with power flow over transmission line k in hour t denoted as e_{kt}. The demand for each hour t at each bus of the network n is denoted as D_{nt}. Minimum and maximum run limits of units are represented as P_g^- and P_g^+ respectively. Operating constrains are denoted compactly in terms of a feasible set $\mathcal{D}$, and vectors are denoted in bold. Thus, the notation $(\mathbf{p}, \mathbf{e}, \mathbf{u}, \mathbf{v}) \in \mathcal{D}$ 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 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 demand) can be described as follows:

$$
\begin{aligned}
&(UC): \min_{p,e,u,v} \sum_{g \in G} \sum_{t \in T} (K_g u_{gt} + S_g v_{gt} + C_g p_{gt}) \\
&s.t. \sum_{g \in G_n} p_{gt} = D_{nt}, n \in N, t \in T \\
&P_g^- u_{gt} \le p_{gt} \le P_g^+ u_{gt}, g \in G, t \in T \\
&e_{kt} = B_k(\theta_{nt} - \theta_{mt}), k = (m,n) \in K, t \in T \\
&(\mathbf{p}, \mathbf{e}, \mathbf{u}, \mathbf{v}) \in \mathcal{D}, \qquad (1)
\end{aligned}
$$

The set of generators located in each bus $n \in N$ is denoted by G_n with $G = \cup_{n \in N} G_n$. The horizon T is 24 hours, with hourly increments. A detailed formulation of the constraints represented by the domain $\mathcal{D}$ can be found in Papavasiliou and Oren [22]. Although the constraint set is not fully presented here, following standard unit commitment models we note that the constraint set can be expressed as a set of linear inequalities on both continuous and integer decision variables.

A. Stochastic Unit Commitment

The stochastic formulation follows the model of Ruiz et al. [6] and involves a two-stage process, where the set of uncertain scenarios is represented as S with each scenario $s \in S$ occurring with a probability π_s. First-stage unit commitment and startup decisions are represented respectively as w and z and apply for those generators 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:

$$
\begin{aligned}
&(SUC): \\
&\min_{p,e,u,v,w,z} \sum_{g \in G} \sum_{s \in S} \sum_{t \in T} \pi_s (K_g u_{gst} + S_g v_{gst} + C_g p_{gst}) \\
&s.t. \sum_{g \in G_n} p_{gst} = D_{nst}, n \in N, s \in S, t \in T \\
&P_{gs}^- u_{gst} \le p_{gst} \le P_{gs}^+ u_{gst}, g \in G, s \in S, t \in T \\
&e_{kst} = B_{ks}(\theta_{nst} - \theta_{mst}), k = (m,n) \in K, s \in T, \\
&t \in T \\
&(\mathbf{p}, \mathbf{e}, \mathbf{u}, \mathbf{v}) \in \mathcal{D}_s, s \in S \\
&u_{gst} = w_{gt}, v_{gst} = z_{gt}, g \in G_s, s \in S, t \in T \qquad (2)
\end{aligned}
$$

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. Scenarios represent the realization of hourly renewable supply production, which results in uncertain net demand D_{nst} in each bus, as well as the loss of generators for the entire day (in which case the capacity limits of a generator are $P_{gs}^- = P_{gs}^+ = 0$, which forces a unit to produce zero output), and the loss of lines (in which case the susceptance of a line is $B_{ks} = 0$, which forces power flow over the line to equal zero).

B. Scenario-Based Security-Constrained Unit Commitment

Note that load shedding is permitted in the stochastic unit commitment model of Eq. (2), with lost load incurring a high penalty in the objective function. Loads $l \in L$ are therefore represented as a dummy generator with second-stage production decisions p_{lst} and a marginal cost equal to the value of lost load. As a result, the feasible region of each scenario, $\mathcal{D}_s$, is non-empty for any choice of first-stage decision variables w_{gt}, z_{gt}.

In a security-constrained model discrete disturbances are accounted for by requiring that the system be capable of withstanding any element failure. This implies that each scenario s now consists of at most a single contingency.

Following the model of Wu et al. [11], we account for continuous disturbances (net demand forecast errors) by associating a renewable supply outcome with each scenario s rather than imposing exogenous reserve requirements. Scenarios that involve no contingency are weighed with a positive probability in the objective function π_s of Eq. (3), whereas scenarios that involve contingencies are only included in the constraint set. The feasible region is equal to $\mathcal{D}_s$ with the additional constraint that $p_{lst} = 0$ for load shedding, and in contrast to (SUC) there may be choices of first-stage decisions for which the model is infeasible (i.e. $\{\mathcal{D}_s, p_{lst} = 0, u_{gst} = w_{gt}, v_{gst} = z_{gt}\} = \emptyset$).

$$
\begin{aligned}
&(SCUC):\\
&\min_{p,e,u,v,w,z} \sum_{g\in G}\sum_{s\in S}\sum_{t\in T}\pi_s(K_g u_{gst} + S_g v_{gst} + C_g p_{gst})\\
&s.t. \sum_{g\in G_n} p_{gst} = D_{nst}, n \in, s \in S, t \in T\\
&P^-_{gs}u_{gst} \le p_{gst} \le P^+_{gs}u_{gst}, g \in G, s \in S, t \in T\\
&e_{kst} = B_{ks}(\theta_{nst} - \theta_{mst}), k = (m,n) \in K, s \in S, t \in T\\
&(\mathbf{p},\mathbf{e},\mathbf{u},\mathbf{v}) \in \mathcal{D}_s, s \in S\\
&p_{lst} = 0, l \in L, s \in S, t \in T\\
&u_{gst} = w_{gt}, v_{gst} = z_{gt}, g \in G, s \in S, t \in T \qquad (3)
\end{aligned}
$$

C. Scenario Selection

The selection of scenarios in the stochastic unit commitment model of Section II-A is based on an idea inspired by importance sampling [22]. A large number of candidate scenarios $\omega \in \Omega$ are evaluated in terms of their cost impact to the system $C_D(\omega)$, where this cost impact is evaluated against an easily computable deterministic unit commitment model. Candidate scenarios are then selected to enter the set of selected scenarios S by sampling according to a probability which is proportional to their cost impact. These scenarios are assigned a probability π_s in the objective function of (SUC) which is inversely proportional to their cost impact $C_D(\omega)$ in order to un-bias their selection. Note that in the stochastic programming formulation, a scenario s can involve any number of contingencies and not necessarily a single contingency.

In the case of the scenario-based security-constrained model, the set S is generated by the Cartesian product of a set of renewable supply outcomes with the no-contingency outcome and the most severe single-element contingencies in the system. The set of scenarios that involve the no-contingency outcome are assigned an equal positive probability in the objective function, whereas the scenarios involving single-element contingencies have no direct impact on the objective function through their weight, $\pi_s = 0$, but only through their presence in the constraint set.

III. Solution Methodology

In the following section we present two decomposition methods for solving (SUC) and $(SCUC)$, as well as distributed implementations of the decomposition algorithms.

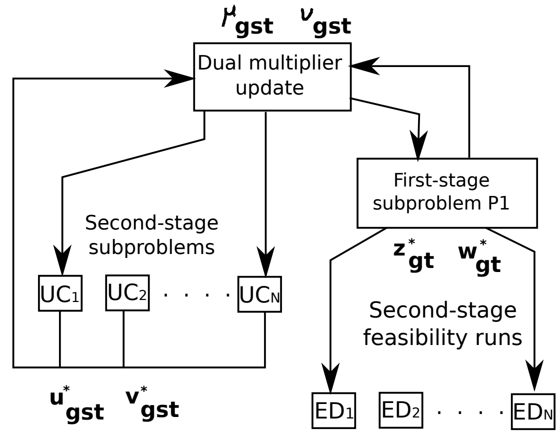

Fig. 1. The parallel implementation of the Lagrangian relaxation algorithm.

A. Lagrangian Relaxation

The Lagrangian relaxation algorithm relies on the observation that the relaxation of the non-anticipativity constraints in (SUC) results in unit commitment subproblems that are independent across scenarios. The Lagrangian dual function is obtained as:

$$
\begin{aligned}
\mathcal{L} = &\sum_{s\in S}\pi_s(\sum_{g\in G}\sum_{t\in T}(K_g u_{gst} + S_g v_{gst} + C_g p_{gst})\\
&+ \sum_{g\in G_s}\sum_{t\in T}(\mu_{gst}(u_{gst} - w_{gt}) + \nu_{gst}(v_{gst} - z_{gt}))\qquad (4)
\end{aligned}
$$

The problem is solved by maximizing the Lagrangian dual function using the sub-gradient algorithm. The solution of the Lagrangian involves one second-stage unit commitment problem for each scenario $(P2_s)$, and one first-stage optimization $(P1)$. The first-stage optimization is formulated as:

$$
(P1): \max_{w,z\in\mathcal{D}_1} \sum_{g\in G_s}\sum_{s\in S}\sum_{t\in T}\pi_s(\mu_{gst}w_{gt} + \nu_{gst}z_{gt}) \qquad (5)
$$

where $\mathcal{D}_1$ represents the minimum up and down time constraints of slow units $g \in G_s$.

The solution of the Lagrangian dual provides a lower bound for the model. By introducing redundant second-stage decision variables on startup decisions, we are able to enforce minimum up and down times on slow units, as in Eq. (5). Given these unit commitment schedules, we can solve an economic dispatch model (ED_s), which is $(P2_s)$ with u_{gst}, v_{gst} fixed for $g \in G_s$. This provides an upper bound that can be used for obtaining feasible solutions at each iteration as well as a duality gap. This duality gap is used as a termination criterion. The algorithm is parallelized both in the solution of $(P2_s)$, as well as the solution of (ED_s), as indicated in Fig. 1. Further details about the solution methodology are discussed in Papavasiliou et al. [5].

B. Benders Decomposition

Security-constrained unit commitment can be approximated as a special case of the (SUC) model presented in Eq. (2) when $\pi_s > 0$ for the no-contingency scenarios, and $\pi_s = 0$ for scenarios involving contingencies. This implies that the constraints associated with each contingency scenario are enforced in the constraint set, but are not weighed in the objective function. This remains an approximation of $(SCUC)$ since the constraint p_{lst}, $l \in L$, is not enforced in (SUC).

In principle, this approximation of the security-constrained unit commitment problem can be solved by using the solution algorithm of Section III-A. In practice this approach presents convergence problems when solved by Lagrangian relaxation. The dual function is not increasing even when the step size is reduced to a very small amount, and the unit commitment schedule of slow generators is inverted after each iteration.

This motivates a Benders decomposition scheme for solving the problem. This can be justified by the fact that all feasibility constraints can be satisfied with only a few feasibility cuts associated with the most severe contingencies in the system. Optimality cuts can be defined by solving only those few scenarios associated with the no-contingency outcome. The advantage of using a Benders decomposition scheme is that the generation of feasibility cuts and optimality cuts can be parallelized, which implies that the second stage of the model is no more the computational bottleneck. The algorithm that we propose in this paper requires two assumptions:

Assumption 1:. In order to maintain the convexity of the second-stage value function, it is necessary to assume that second-stage problems are continuous. Therefore, we impose the assumption that unit commitment decisions have to be fixed for all generators in the network from the first stage. This is contrasted to the Lagrangian relaxation algorithm that can involve integer decisions in the second stage for fast generators $g \in G_f = G - G_s$.

Assumption 2: The generation of feasibility cuts according to Van-Slyke and Wets [23] removes one candidate integer solution at each iteration, however this process can easily stall when there is a large number of candidate integer solution combinations that need to be tested before a feasible solution can be obtained, as is the case in the stochastic unit commitment problem. By assuming away ramping constraints in $(SCUC)$, we obtain a feasible region $(\mathcal{D}_{st})$ that is decomposable both by time period as well as scenario. Rather than using the feasibility cuts of Van-Slyke and Wets [23], we then impose the constraints represented by $\mathcal{D}_{st}$ in the first-stage problem for the scenario and time period that represents the most severe contingency given the current candidate integer solution. The motivation is that accounting for the most severe contingency in the first stage of the problem is capable of satisfying most operating constraints associated with less severe contingencies. The algorithm is presented in Fig. 2.

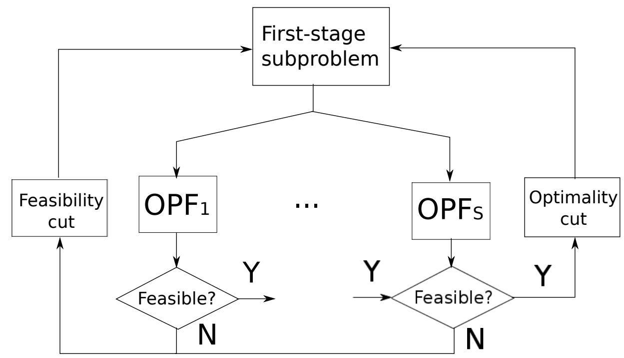

Fig. 2. The parallel implementation of the Benders decomposition algorithm.

IV. 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. There are 82 fast units, with a total capacity of 9,156.1 MW and 42 slow units with a total capacity of 19,225.4 MW. The fuel mix of the generators and a schematic are presented in [5]. The value of lost load is assumed equal to 5,000 \$/MWh.

The wind penetration level that we analyze corresponds to the 2030 wind integration targets of California for a typical spring weekday. The wind model is calibrated against one year of data from the National Renewable Energy Laboratory. The wind power production time series model is described in detail by Papavasiliou and Oren [24]. The model captures temporal correlations of wind speed, the nonlinear conversion of wind speed to wind power, the locational correlations of the wind sites under consideration, as well as systematic seasonal and diurnal characteristics of the data set. The wind power production time series model was used both in order to generate scenarios for the unit commitment optimization models, as well as for generating outcomes for the Monte Carlo simulation of the performance of the two different unit commitment policies. The set of outcomes that were used for the Monte Carlo performance evaluation were different from the scenarios that were used as input to the unit commitment models.

Both formulations were solved for 30 scenarios in order to compare the two models on a fair basis. The scenario selection algorithm of Section II-C was used for selecting and weighing the scenarios of (SUC). The input for the $(SCUC)$ model was generated from the Cartesian product of ten wind power production scenarios with the no-contingency case as well as the two most severe contingencies in the network, namely the failure of each of the two nuclear units in the network. Both algorithms were implemented in the Java callable library of CPLEX 12.4, and parallelized using the Message Passing Interface (MPI). The code was

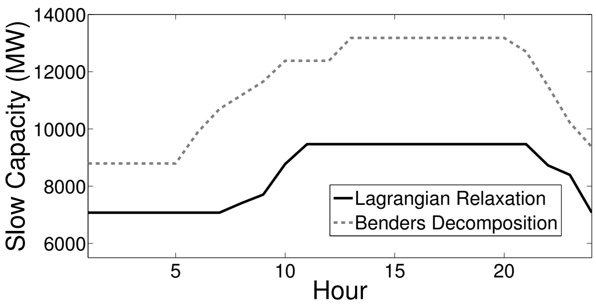

Fig. 3. Hourly day-ahead capacity committed by each model.

implemented on a high performance computing cluster in the Lawrence Livermore National Laboratory, with 8 CPUs per node, 2.4 GHz and 10 GB per node.

Each unit commitment policy was evaluated against 1,000 Monte Carlo outcomes of wind power production and contingencies. We assume a probability of generator failure equal to 1 % [25] and a probability of transmission line failure equal to 0.1 % [26].

For the implementation of the (SUC) algorithm, $(P1)$ and $(P2_s)$ were run for 80 iterations. For the last 40 iterations, (ED_s) was run for each scenario in order to obtain a feasible solution and an upper bound to the problem. The best upper bound was evaluated at 5.911 $M, while the best lower bound was evaluated at 5.868 $M.

The Benders decomposition algorithm required 31 iterations to converge. During these iterations, either feasibility cuts were added to the first-stage program, or a new approximation of the value function was generated, along with an estimate of the gap in the current candidate unit commitment solution. The first feasible unit commitment schedule was detected in iteration 19. The gap at iteration 19, was equal to 5.855 $M. Subsequently, the value function approximation improved around the neighborhood of the optimal solution, and although 4 more feasibility cuts were added in the remaining iterations, the algorithm eventually terminated after 31 iterations.

A. Relative Peformance

The hourly day-ahead capacity committed by each model in each hour of the day is shown in Fig. 3. We note that the $(SCUC)$ model is committing significantly more capacity than the (SUC) model. This can be attributed to Assumption 2 of Section III-B. Due to the fact that the Benders decomposition algorithm requires that all units be committed in the day-ahead time frame, the resulting policy is quite conservative. The cost performance of the two approaches in the Mote Carlo simulation is shown in Table I. We note that the (SUC) model outperforms the $(SCUC)$ model by 5.4 % relative to the average daily cost of the $(SCUC)$ model in terms of expected cost performance.

B. Running Time

The running time of the Benders decomposition algorithm is shown in Fig. 4. The speedup of the algorithm is due to the parallelization of the continuous DC optimal power flow problems that are required for generating feasibility and optimality cuts (see Fig. 2). The marginal benefits vanish beyond 15 processors. The entire model requires 26.6 minutes to solve in a fully serial implementation, versus 14.8 minutes in a fully parallel implementation. We note that the benefits of parallelism are expected to increase as we increase the number of contingencies or wind scenarios considered in the model. However, as an excessive number of second-stage problems is added to the $(SCUC)$ model, additional feasibility cuts are required in order to generate feasible unit commitment schedules. This may result in a non-decomposable first-stage problem that is excessively large, and for which distributed computation can offer no speedup benefits. In that case, the first-stage problem will dominate the total running time of the problem. We have encountered this behavior in an instance of the $(SCUC)$ problem with 100 contingencies (which results in 1,000 scenarios when the contingencies are interleaved with 10 wind scenarios), and in future research we intend to explore alternative approaches for solving larger instances of the problem.

TABLE I
DAILY OPERATING COST ($1,000)

	Startup	Min. load	Load shed	Fuel	Total
Benders	66.5	1,205.3	0	4,687.3	5,959.1
LR	106.0	699.4	0.3	4,831.5	5,637.2

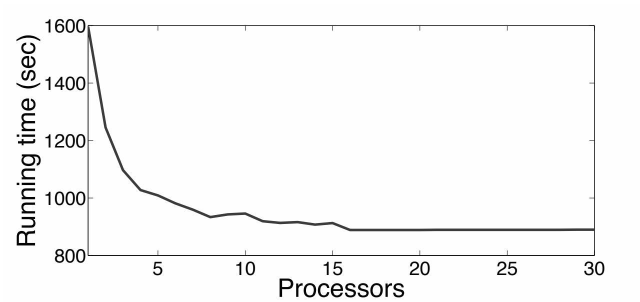

Fig. 4. The running time of the Benders decomposition algorithm as a function of processor number.

The running time of the Lagrangian relaxation algorithm is shown in Fig. 5. The marginal benefits of parallelization vanish beyond 15 processors. The solution time of the Lagrangian relaxation algorithm ranges between 15.8 hours for the fully serial implementation to 47.7 minutes in the fully parallel implementation. The benefits of parallelization are evident in this example, as they enable us to reduce the solving time of the original problem to a time horizon that is acceptable for operational purposes. In contrast to the $(SCUC)$ model, the proposed Lagrangian relaxation can scale to a very large number of scenarios provided that a sufficient number of processors is available.

V. Conclusions

We compared a stochastic unit commitment model and a scenario-based security constrained unit commitment model

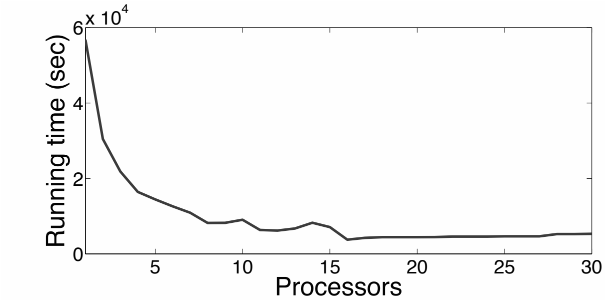

Fig. 5. The running time of the Lagrangian relaxation algorithm as a function of processor number.

for solving the unit commitment problem. We tested the two approaches on a test case of the California ISO interconnected with the Western Electricity Coordinating Council. We observed that the security-constrained model commits significantly greater quantities of day-ahead capacity and outperforms the stochastic unit commitment model in terms of load shedding. Instead, the stochastic unit commitment model outperforms the security-constrained model in terms of expected cost by reducing minimum load, startup and fuel costs. Both decomposition algorithms used for solving each model benefit from parallelization, although running time in the Benders algorithm is dictated by the first-stage subproblem. Further research is required in order to solve larger instances of the security-constrained model.

ACKNOWLEDGMENT

The work described in this paper was made possible by funding provided by the U.S. Department of Energy for "The Future Grid to Enable Sustainable Energy Systems", an initiative of the Power Systems Engineering Research Center, and by an LDRD grant from the Lawrence Livermore National Laboratory.

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. M. Morales, A. J. Conejo, and J. Perez-Ruiz, "Economic valuation of reserves in power systems with high penetration of wind power," *IEEE Transactions on Power Systems*, vol. 24, no. 2, pp. 900–910, May 2009.

[5] 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.

[6] P. A. Ruiz, R. C. Philbrick, E. Zack, K. W. Cheung, and P. W. Sauer, "Uncertainty management in the unit commitment problem," *IEEE Transactions on Power Systems*, vol. 24, no. 2, pp. 642–1651, May 2009.

[7] T. Shiina and J. R. Birge, "Stochastic unit commitment problem," *International Transactions on Operations Research*, vol. 11, no. 95, pp. 19–32, 2004.

[8] 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.

[9] 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.

[10] 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.

[11] 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.

[12] R. Jiang, M. Zhang, G. Li, and Y. Guan, "Robust unit commitment with wind power and pumped storage hydro," *IEEE Transactions on Power Systems*, vol. 27, no. 2, pp. 800–810, 2012.

[13] 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.

[14] 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.

[15] M. Vrakopoulou, K. Margellos, J. Lygeros, and G. Andersson, "Probabilistic guarantees for the N-1 security of systems with wind power generation," in *Probabilistic Methods Applied to Power Systems*, Istanbul, Turkey, June 2012, pp. 858–863.

[16] D. Falcao, "High performance computing in power system applications," *Lecture Notes in Computer Science*, vol. 1215, pp. 1–23, 1997.

[17] 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.

[18] 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.

[19] 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.

[20] 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.

[21] 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.

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

[23] R. M. V. Slyke and R. Wets, "L-shaped linear programs with applications to optimal control and stochastic programming," *SIAM Journal on Applied Mathematics*, vol. 17, no. 4, pp. 638–663, July 1969.

[24] A. Papavasiliou and S. S. Oren, "Stochastic modeling of multi-area wind production," in *12th International Conference on Probabilistic Methods Applied to Power Systems*, June 2012.

[25] 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.

[26] 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.