

Application of Stochastic Dual Dynamic Programming to the Real-Time Dispatch of Storage Under Renewable Supply Uncertainty

Anthony Papavasiliou , Member, IEEE, Yuting Mou, Student Member, IEEE, Léopold Cambier, and Damien Scieur

Abstract—This paper presents a multistage stochastic programming formulation of a transmission-constrained economic dispatch subject to multiarea renewable production uncertainty, with a focus on optimizing the dispatch of storage in real-time operations. This problem is resolved using stochastic dual dynamic programming. The applicability of the proposed approach is demonstrated on a realistic case study of the German power system calibrated against the solar and wind power integration levels of 2013–2014, with a 24-h horizon and 15-min time step. The value of the stochastic solution relative to the cost of a deterministic policy amounts to 1.1%, while the value of perfect foresight relative to the cost of the stochastic programming policy amounts to 0.8%. The relative performance of various alternative real-time dispatch policies is analyzed, and the sensitivity of the results is explored.

Index Terms—Dynamic programming, economic dispatch, renewable energy, stochastic programming, storage.

I. INTRODUCTION

DAY-AHEAD and real-time power system operations have become increasingly intricate in recent years due to the large-scale integration of renewable energy resources. Flexible resources can be used in order to mitigate the operational challenges of renewable energy supply. This paper focuses on the role of storage, specifically pumped hydro resources, in real-time operations.

Current industry practice for conducting real-time operations relies on economic dispatch with a relatively limited look-ahead horizon and a deterministic forecast of renewable energy supply, which is fairly accurate given the limited look-ahead of existing dispatch models. In systems with significant amounts of storage and renewable resources, this practice may lead to

inefficiencies [1]. The deployment of a wide variety of storage resources, including pumped hydro storage, utility-scale batteries, flywheels, deferrable demand including electric vehicle storage, and distributed residential storage, requires a real-time dispatch procedure which is capable of long-range look-ahead and potent uncertainty management. Stochastic programming is suitable for this task. This paper proposes a real-time economic dispatch procedure based on multi-stage stochastic optimization which can be used for the management of storage resources.

The beneficial role of storage in mitigating renewable supply variability and uncertainty in short-term (day-ahead and real-time) operations has been studied extensively in the literature. Early work on the economic implications of optimizing storage in short-term operations includes Swider *et al.* [2] and Tuohy *et al.* [3]. The coexistence of renewable supply uncertainty and storage in short-term operations leads naturally to the consideration of stochastic programming models, which are well suited for optimizing under uncertainty in problems over multiple time periods with significant coupling of operational decisions over time. This is in direct analogy to the long tradition of stochastic programming in resolving medium-term planning models of hydroelectric storage under precipitation uncertainty, which was pioneered by the early work of Pereira and Pinto [4].

The majority of the literature on the optimal deployment of storage in short-term operations relies on two-stage stochastic unit commitment formulations. Khodayar *et al.* [5] consider a two-stage scheduling problem where the first stage corresponds to day-ahead unit commitment and the second stage involves optimal deployment of storage against realizations of load forecast errors, wind forecast errors and line and generator outages. Ponzo *et al.* [6] also consider a two-stage stochastic unit commitment model that determines day-ahead commitment and reserve decisions, where the system is deployed in real time against realized load and wind forecast errors. Deane *et al.* [7] use stochastic programming in order to derive weekly and daily reservoir targets for pumped hydro resources, where reservoir levels are considered as non-anticipative decisions. Li *et al.* [8] consider a stochastic unit commitment model where conventional units are dispatched in the first stage, and scenario-dependent unit commitment decisions are determined in the second stage. Nürnberg and Römis [9] consider two-stage stochastic unit commitment with wind uncertainty and fuel cost uncertainty, and focus on the

Manuscript received March 8, 2017; revised June 14, 2017 and August 1, 2017; accepted August 30, 2017. Date of publication September 1, 2017; date of current version March 20, 2018. This work was supported in part by the ENGIE Chair in Energy Economics and Energy Risk Management and in part by the ENGIE-Electrabel ColorPower Grant. Paper no. TSTE-00224-2017. (Corresponding author: Anthony Papavasiliou.)

A. Papavasiliou and Y. Mou are with the Center for Operations Research and Econometrics, Université Catholique de Louvain, Louvain-la-Neuve 1348, Belgium (e-mail: anthony.papavasiliou@uclouvain.be; yuting.mou@uclouvain.be).

L. Cambier is with the Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA 94305 USA (e-mail: lcambier@stanford.edu).

D. Scieur is with the École Normale Supérieure, Paris 75005, France (e-mail: damien.scieur@gmail.com).

Digital Object Identifier 10.1109/TSTE.2017.2748463

hourly dispatch of pumped hydro resources over the duration of a week.

Alternative approaches to managing uncertainty through the optimal short-term dispatch of storage are also considered in the literature. Jiang *et al.* [10] use a two-stage robust optimization model for the optimal management of storage in order to address wind power uncertainty. Wen *et al.* [11] consider a unit commitment model for using utility-scale storage in order to determine post-contingency corrective actions.

In contrast to the previous literature that has been dedicated to day-ahead unit commitment models for optimizing the operation of storage under renewable supply uncertainty, the literature on multi-stage economic dispatch under uncertainty in real-time operations is relatively less developed. Lorca and Sun [12] analyze two-stage robust optimization models for economic dispatch, and extend this work further in Lorca *et al.* [13]. As in the case of the present paper, this work is motivated by the growing importance of multi-stage real-time economic dispatch under uncertainty due to the variability and uncertainty of wind power output. Wang and Hobbs [14], [15] have also investigated this problem recently, with a focus on analyzing flexible ramping products and the extent to which they can approach the ideal outcome of real-time stochastic optimization. Safta *et al.* [16] also consider multi-period stochastic economic dispatch from the point of view of developing efficient sampling methods of the underlying uncertainty. Phan and Ghosh [17] develop decomposition algorithms for a two-stage stochastic optimal power flow model, where non-linear power flow constraints are considered explicitly in the model.

An alternative approach to the management of real-time uncertainty which has been analyzed in the literature is the optimization of participation factors of flexible resources in automatic generation control (AGC) in order to alleviate load and renewable supply uncertainty, while satisfying probabilistic constraints on the operation of the network. Jabr [18] introduces adjustable robust optimal power flow, and presents a model in both a static as well as a dynamic setting. Bienstock *et al.* [19] analyze real-time optimal power flow under probabilistic constraints. Lubin *et al.* [20] extend these models by developing a solution procedure that accounts for distributionally robust optimization. Convex formulations of chance constraints for the case of affine control of AGC are provided by Summers *et al.* [21]. Jabr *et al.* [22] optimize the utilization of storage-type resources using affinely adjustable AGC control.

In the backdrop of the aforementioned literature, the present paper proposes a multi-stage stochastic programming formulation of economic dispatch, with a specific focus on optimizing the management of storage resources in real-time operations. A major drawback of the previously cited day-ahead models is the assumption that uncertainty about the intra-day evolution of renewable supply uncertainty is revealed at once in the beginning of the second stage. The alternative of considering a multi-stage stochastic programming model results in computational challenges that are overcome in the present paper through the use of stochastic dual dynamic programming (SDDP).

The contributions of this paper can be summarized as follows: (i) The SDDP algorithm has been most successfully

applied in the context of medium-term multi-stage hydrothermal scheduling under rainfall uncertainty for handling water levels of hydro reservoirs [4]. This paper demonstrates its computational viability for an application of increasing significance: short-term multi-stage economic dispatch under renewable supply uncertainty for dispatching storage resources. The analogies of the two problems are striking, however the applicability of SDDP in the context of multi-stage stochastic economic dispatch has not been explored in the literature. (ii) The potential value of scheduling storage in real-time operations using value functions obtained from multi-stage stochastic programming is analyzed. The majority of the existing literature focuses on analyzing the value of storage using two-stage day-ahead unit commitment models which are overly optimistic, or economic dispatch models of small scale or limited look-ahead, whereas storage creates temporal coupling over the entire day. The sensitivity of these findings with respect to transmission constraints, ramp constraints, and optimization look-ahead is analyzed. (iii) Practical implementation challenges such that multi-stage stochastic economic dispatch can be embedded in existing power system and power market operations are highlighted, and solutions towards overcoming these challenges are investigated.

The remainder of the paper is organized as follows. Section II provides an introduction to multi-stage stochastic programming and the SDDP algorithm. Section III casts the real-time pumped hydro scheduling problem as a multi-stage stochastic programming problem which can be tackled by SDDP. Section IV presents a case study focused on the German power system, while Section V discusses aspects related to practical implementation. Conclusions are drawn and future lines of research are delineated in Section VI.

II. MULTI-STAGE STOCHASTIC LINEAR PROGRAMMING AND SDDP

A. Multi-Stage Stochastic Linear Programming

This paper considers multi-stage stochastic linear programming problems with discrete time steps $t \in T = \{1, \dots, H\}$, where H is the horizon of the problem. Consider a discrete set of realizations of uncertainty at each time step. Denote Ω_t as the discrete set of outcomes in stage t , and $\Omega_{[t]}$ as the set of possible histories up to stage t . Each $\omega_{[t]} \in \Omega_{[t]}$ has a unique ancestor $A(\omega_{[t]}) \in \Omega_{[t-1]}$. The notation is explained further in Section A of the appendix.

In what follows, the standard assumption of serial independence in SDDP [23] is adopted. It is assumed that at each stage the objective function coefficients c_t and constraint coefficients W_t are deterministic parameters, whereas the right-hand side parameters h_{t,ω_t} and the constraint coefficients T_{t,ω_t} are random. Define decision variables that depend on the history of realizations up to stage t , $x_{t,\omega_{[t]}}$, where $\omega_{[t]} \in \Omega_{[t]}$. Serial independence implies that it is possible to define, for each stage t , a probability measure p_{t,ω_t} for each outcome $\omega_t \in \Omega_t$, from which one can recursively compute $p_{t,\omega_{[t]}} = p_{t-1,A(\omega_{[t]})} \cdot p_{t,\omega_t}$ with $p_{1,\omega_1} = p_{1,\omega_{[1]}} = 1$. The following multi-stage stochastic

linear program in *extended form* can then be defined:

$$\begin{aligned} \min_x \quad & \sum_{t=1}^H \sum_{\omega_{[t]} \in \Omega_{[t]}} p_{t,\omega_{[t]}} c_t^T x_{t,\omega_{[t]}} \\ T_{t,\omega_t} x_{t-1,A(\omega_{[t]})} + W_t x_{t,\omega_{[t]}} &= h_{t,\omega_t}, t \in T, \omega_{[t]} \in \Omega_{[t]} \\ x_{t,\omega_{[t]}} &\geq 0, t \in T, \omega_{[t]} \in \Omega_{[t]} \end{aligned}$$

B. Nested L-Shaped Decomposition Subproblem

The SDDP algorithm decomposes the above problem, which is intractable in extended form, to a collection of subproblems, one for each time stage $t \in T$ and for each outcome $k \in \Omega_t$. Each of these subproblems is referred to as a *nested L-shaped decomposition subproblem (NLDS)* [23]. Exploiting the fact that the value functions of multi-stage stochastic linear programs are piecewise affine functions, the NLDS for stage t and outcome k can be expressed as

$$NLDS_{t,k} : \min_x c_t^T x + \tilde{V}_t(x) \quad (1)$$

$$W_t x = h_{t,k} - T_{t,k} \hat{x}_{t-1} \quad (2)$$

$$x \geq 0 \quad (3)$$

where $\tilde{V}_t(x)$ is a piecewise affine convex function of x which approximates the expected cost-to-go as a function of the decisions x made in stage t and for outcome k . Since $\tilde{V}_t(x)$ is a piecewise affine convex function, $NLDS_{t,k}$ is a linear program that can be solved efficiently. Note that $\hat{x}_{t-1}$ is a trial decision made in stage $t-1$, which is an input, not a decision variable, for $NLDS_{t,k}$. Feasibility cuts are ignored, since the economic dispatch problem that is presented in the next section includes load shedding decision variables which imply that NLDS is always feasible for any x .

C. Lattice

Uncertainty is represented in the present implementation of the SDDP algorithm by a lattice, which is a graphical description of a discrete Markov process¹. Each stage t consists of a set of nodes Ω_t . Each node $k \in \Omega_t$ corresponds to a realization $h_{t,k}$ of the random vector h_{t,ω_t} and a realization $T_{t,k}$ of the random matrix T_{t,ω_t} . Since serial independence is assumed, the specification of the uncertainty model can be completed by associating with each node (t,k) a probability of occurrence $p_{t,k}$, which is independent of the history of realizations up to stage $t-1$.

The linear program $NLDS_{t,k}$ defined in (1)–(3) is stored at each node (t,k) of the lattice. Serial independence implies that each node $k \in \Omega_t$ stores an identical copy of the value function $\tilde{V}(x)$ of (1). In general, the MATLAB toolbox that is presented in the following section can also handle random

objective function coefficients c_t and random constraint matrices W_t . The lattice can generally represent Markov processes whereby the user defines the realization of the random parameter² $\xi_t^T = (c_t^T, h_t^T, \text{vec}(T_t), \text{vec}(W_t))$, and the transition probability from any node in stage $t-1$ to any node in stage t .

The lattice is the core of the FAST toolbox presented in the next section. It stores the description of the uncertainty model, and can be used for simulating, not only the SDDP policy described by $NLDS_{t,k}$, but any heuristic policy that can be described by a linear program.

D. FAST Toolbox

FAST is a MATLAB open-source toolbox developed by the authors which can be used for solving multi-stage stochastic linear programming problems with SDDP. The toolbox can also be used for simulating perfect foresight policies, as well as heuristic policies that can be expressed through linear programming. The toolbox is described in detail in the following link: <https://web.stanford.edu/~lcambier/fast/>. This section briefly covers the features that are relevant to the model implemented in the paper.

The definition of a stochastic program in FAST requires (i) the description of the $NLDS$ (which can be done as in a high-level mathematical programming language, as opposed to manually entering the constraint matrix, objective function coefficients, and right-hand side parameters of the linear program), (ii) the description of the lattice, and (iii) the description of user-defined settings related to the behavior and convergence of the algorithm. The $NLDS$ and the lattice are described in Sections II-B and II-C respectively, this section describes the settings that control convergence.

The SDDP algorithm proceeds by iterating between forward and backward passes. The role of forward passes is to generate bounds and trial decisions, while backward passes generate optimality cuts that approximate the value function at each stage. The algorithm is described in detail by Pereira and Pinto [4]. The toolbox is compatible with a variety of commercial solvers, including CPLEX and Gurobi.

The forward pass proceeds by generating a user-defined number of Monte Carlo samples, M . Each Monte Carlo sample generates a realization of the random parameters h_{t,ω_t} and T_{t,ω_t} over the entire horizon of the problem. This results in a trajectory of trial decisions, $\hat{x}^{(i)} = (\hat{x}_1^{(i)}, \dots, \hat{x}_H^{(i)})$ (which are obtained by solving the $NLDS$ of each stage t and the realization of the Monte Carlo sample) as well as an associated sample cost over the entire horizon, z_i (since the sequence of decisions $\hat{x}^{(i)}$ is a feasible, not necessarily optimal, reaction to the realized sample of uncertainty). A statistical upper bound can then be obtained by the sample average $\bar{z} = \frac{\sum_{i=1}^M z_i}{M}$. Within FAST, the user can periodically run a large number of forward passes in order to obtain an accurate estimate of the upper bound.

¹This assumption, although restrictive, can accommodate a large class of practically useful discrete time stochastic processes, including autoregressive moving average models [24]. On the upside, this assumption yields computational gains by allowing the sharing of dual multipliers for the generation of optimality cuts. This is illustrated in Fig. 7 in the appendix.

²The operator $\text{vec}(\cdot)$ applied to a matrix expands the columns of the matrix to a vector, i.e., $\text{vec}(A) = (A_{:,1}^T, \dots, A_{:,n}^T)$ where n is the number of columns of A and $A_{:,j}$ is the j -th column of the matrix.

The backward pass generates optimality cuts that approximate the value function $\hat{V}_t(x)$ around the trial points $\hat{x}_t^{(i)}$ generated by the forward pass. These approximations (which can be shown to be global under-estimators of the true value function) are generated by solving each *NLDS* and exploiting the dual optimal multipliers, similarly to the standard L-shaped method [25]. Since the approximations $\hat{V}_t(\cdot)$ are under-approximations of the true value functions, the solution of *NLDS*₁ furnishes a lower bound $\underline{z}$. The optimality cuts generated after the algorithm has converged can be used for simulating the resulting policy. FAST can also prune optimality cuts in order to keep the size of each *NLDS* tractable.

Various convergence criteria can be selected within FAST, including a minimum and maximum number of iterations, the classical convergence criterion³ of Pereira and Pinto [4], as well as a small standard deviation criterion according to which the algorithm terminates only when the standard deviation of the mean cost $\bar{z}$ is sufficiently low, relative to the lower bound $\underline{z}$. The convergence criteria of SDDP are discussed in detail by De-Mello *et al.* [27].

III. MODEL DESCRIPTION

Since the focus of the present paper is on real-time dispatch, unit commitment decisions are assumed to be fixed. The economic dispatch model presented in this section specifically focuses on the management of pumped hydro resources throughout the day.

A. Renewable Supply Model

Despite the assumption of serial independence, the multi-stage stochastic linear programming problem described in Section II-A can be used to describe a rich family of stochastic processes [24], [28]. In order to motivate the approach adopted in this paper, consider first how a first-order autoregressive process AR(1) would be modeled according to constraints (2). This paper focuses in modeling renewable supply in the presence of forecast errors, which can be expressed as follows:

$$re_t = c + \phi \cdot re_{t-1} + \epsilon_t,$$

$$p_t = RF_t + re_t$$

where c is the constant term of the AR(1) model, ϕ is the autoregressive coefficient, ϵ_t are independent identically distributed normal variables, re_t is the forecast error which is modeled as an AR(1) process, p_t is the renewable energy production in stage t , and RF_t is the forecast production in stage t . The idea is to include the forecast error re in the state vector x of (2), and to model the AR noise ϵ_t as the serially independent stochastic input h_{t,ω_t} of (2). Whereas it might seem more natural to model

³The convergence criterion of Pereira and Pinto [4] stipulates that the algorithm should terminate when the lower bound $\underline{z}$ lies within the interval

$[\bar{z} - 2\sigma, \bar{z} + 2\sigma]$, where $\sigma = \sqrt{\frac{\sum_{i=1}^M (z_i - \bar{z})^2}{M^2}}$ is an estimate of the standard deviation of the sample average cost. Effectively, the termination criterion requires that the lower bound lie within the 95% confidence interval of the upper bound, however this may lead to premature convergence [26].




Fig. 1. The scatter plot of wind and solar production day-ahead forecasts versus historical forecast error in Germany for 2013 and 2014.

forecast error directly as the random parameter, instead of representing it as a state variable, in this way it is possible to represent autoregressive processes while respecting serial independence.

There are two drawbacks with this approach, which motivate the alternative modeling approach adopted in this paper. The first drawback is that renewable production becomes negative⁴. The second drawback is that the forecast error under an AR(1) model would have constant variance, whereas experimental evidence suggests that the variance of forecast errors increases as the renewable forecasts increase, as shown in Fig. 1.

Cabral [29] recently proposed a multiplicative autoregressive model for representing inflow uncertainty in hydrothermal planning models, which accounts for temporal correlations in the inflow process, while maintaining the non-negativity of the process. This overcomes the first drawback mentioned in the previous paragraph. In order to overcome the second drawback, this paper proposes that the multiplicative autoregressive model of Cabral be used in order to model the *ratio* of forecast error to renewable supply forecast, instead of the renewable supply process.

Denote the ratio of renewable production realizations and day-ahead forecasts as $y_t = p_t / RF_t$. Then the autoregressive multiplicative model of Cabral can be used for modeling the ratio y_t as follows:

$$y_{t+1} = (c + \phi \cdot y_t) \cdot \eta_t \quad (4)$$

$$p_t = RF_t \cdot y_t \quad (5)$$

where η_t are independent, identically distributed.

In order to generate a lattice for SDDP, the random parameter η_t is discretized at every time stage. In this paper, the discretization is based on discretizing the normal distribution into ten

⁴Alternatively, one could impose non-negativity on p_t , however then one would have to include slack variables with high penalties in order to ensure that the error re remains in a range such that $p_t \geq 0$, which would interfere with the optimality cuts generated by the SDDP algorithm.




Fig. 2. Histogram of historical renewable supply observations (left panel) and modeled renewable supply (right panel).

evenly spaced quantiles:

$$\eta_{t,k} = \mu_t + \sigma_t \cdot \Phi^{-1} \left(\frac{k}{10} - 0.05 \right)$$

where μ_t and σ_t are the mean and standard deviation of the estimated noise in stage t respectively, Φ is the cumulative distribution function of the standard normal distribution, and $\eta_{t,k}$ is the value of the noise in node k of stage t . Fig. 2 compares the probability density function obtained from historical data with the probability density function produced by the proposed stochastic model. The quality of the model is further validated by comparing the energy score⁵ of the multiplicative first-order autoregressive model to the energy score of an additive first-order autoregressive model over 100 days of data. The multiplicative model employed in this paper achieves an average energy score of 0.26, and outperforms the additive model, the average energy score of which amounts to 0.46.

B. Stochastic Multi-Period Economic Dispatch

This section formulates the multi-stage stochastic economic dispatch problem, by describing the *NLDS*, i.e., the objective function (1) and constraints (2) for a given stage $t \in T$ and history $\omega_{[t]} \in \Omega_{[t]}$. The t and $\omega_{[t]}$ indices are dropped in order to keep the notation concise. The notation is described in detail in Section VI-B of the appendix.

The objective is to minimize the expected cost of power production and load shedding:

$$\frac{1}{4} \left(\sum_{n \in N} VOLL \cdot ls_n + \sum_{g \in G} c_g \right)$$

where G is the set of generators, N is the set of buses, ls_n is the load shedding in bus n , and c_g is the cost of generator g . The 0.25 coefficient in the objective function is due to the assumed time resolution of 15 minutes for each time step.

Production cost is defined as a piecewise affine function:

$$c_g \geq FC_g \cdot (A_{g,m} \cdot U_g + B_{g,m} \cdot p_g), g \in G, m = 1, \dots, 3$$

where p_g is the power production of a generator, FC_g is the fuel cost of a generator, $A_{g,m}$ and $B_{g,m}$ are the intercept and slope that describe segment m of the heat rate curve of a generator, and

⁵The energy score is a proper (i.e., a perfect forecast will result in the best score), negatively-oriented (i.e., lower is better) score that quantifies both the skill (accuracy) and sharpness (spread) of a scenario set [30].

U_g is a binary variable indicating whether a generator has been committed or not. Note that U_g is determined in the day-ahead time frame, and is a fixed parameter in the real-time dispatch.

The transmission network is represented as a directed graph (N, K) , where K is the set of transmission lines. At each node of the network, the following power balance constraint is imposed:

$$\begin{aligned} \sum_{l \in L_n} D_l + \sum_{g \in PH_n} pd_g + \sum_{k \in (n, \cdot)} f_k + \sum_{n \in N} ps_n = \sum_{g \in G_n} p_g \\ + \sum_{g \in PH_n} pp_g + \sum_{g \in GR_n} RF_g \cdot y_g + ls_n + \sum_{k \in (\cdot, n)} f_k, n \in N \end{aligned}$$

where L_n is the set of loads in location n , PH_n is the set of pumped hydro units in location n , G_n is the set of generators in location n , GR_n is the set of renewable resources in location n , $(n, \cdot)$ is the set of lines whose origin is n , and $(\cdot, n)$ is the set of lines whose destination is n , D_l is the load of consumer l , pd_g is the pumping demand of pumped hydro unit g , pp_g is the pumping supply of pumped hydro unit g , f_k is the flow on line k , ps_n is amount of production shedding in node n , RF_g is the renewable energy forecast of renewable generator g , and y_g is the ratio of renewable production to renewable forecast error for renewable generator g .

Power flow equations⁶ are expressed using a linearized loss-less model of the transmission network:

$$\theta_{\text{hub}} = 0$$

$$f_k = B_k \cdot (\theta_m - \theta_n), k = (m, n) \in K$$

where θ_n is the angle of bus n (with the hub node assumed to be the bus angle reference node) and B_k is the susceptance of line k .

Power flows are limited as follows:

$$-TC_k \leq f_k \leq TC_k, k \in K$$

where TC_k is the flow limit of line k .

The dynamics of pumped hydro units are expressed as follows:

$$s_g = s_{g,t-1} + 0.25 \cdot \left(\Gamma_g^d \cdot pd_g - \frac{pp_g}{\Gamma_g^p} \right), g \in PH$$

where s_g is the stored energy of pumped hydro unit g , Γ_g^d is the pumping efficiency of unit g , and Γ_g^p is the production efficiency of unit g .

The capacity constraints of conventional and pumped hydro units are expressed as follows:

$$pd_g \leq DMax_g, g \in PH$$

$$pp_g \leq PMax_g, g \in PH$$

$$PMin \cdot U_g \leq p_g \leq PMax_g \cdot U_g, g \in G$$

⁶Although mathematically equivalent to a formulation which uses power transfer distribution factors (PTDFs), the bus angle formulation employed in this paper appears to achieve faster convergence for the case study described in Section IV. In particular, whereas the SDDP algorithm with a bus angle formulation converges in 4.3 hours, the same algorithm requires 16.4 hours to converge with a PTDF formulation. This may be due to the fact that the PTDF formulation involves a dense constraint matrix, which can impact the speed of matrix inversion when the *NLDS* are being resolved in the course of the SDDP algorithm.

where $DMax_g$ is the pumping capacity of pumped hydro unit g , $PMax_g$ is the production capacity of pumped hydro or conventional generator g , and $PMin_g$ is the technical minimum of conventional generator g .

The ramp rate constraints of conventional units are expressed as follows:

$$p_g - p_{g,t-1} \leq RU_g \cdot U_g + MTL_g \cdot (1 - U_{g,t-1}), g \in G$$

$$p_{g,t-1} - p_g \leq RD_g \cdot U_g + MTL_g \cdot (1 - U_g), g \in G$$

where RU_g is the 15-minute ramp-up rate of generator g , RD_g is the 15-minute ramp-down rate, and MTL_g is the maximum transition limit of generator g . The maximum transition limit is the power output of a generator after being turned on, or before being turned off.

The following constraints impose bounds on load and production shedding:

$$ls_n \leq \sum_{l \in L_n} D_l, n \in N$$

$$ps_n \leq \sum_{g \in G_n} p_g + \sum_{g \in GR_n} RF_g \cdot y_g, n \in N$$

The dynamics of forecast errors can be described by (4). Finally, the following non-negativity constraints are required:

$$ls_n, ps_n, s_g, pd_g, pp_g \geq 0, n \in N, g \in PS$$

IV. CASE STUDY

This section presents a case study of the German system. Germany is chosen as the most relevant example within Europe, due to its leading position in the integration of wind and solar power. The rapid growth of renewable energy supply in Germany is evidenced by the evolution of installed capacity in recent years. The installed capacity of solar power evolved from 36.3 GW in 2013 to 39.8 GW in 2015, while the installed capacity of wind power evolved from 34.0 GW in 2013 to 44.5 GW in 2015 [source: German Federal Ministry of Economic Affairs and Energy]. The ratio of historically observed renewable energy production to total energy production for the interval of the data which was used in the study (January 2013-December 2014) is 18.7%. The installed capacity of conventional generators, which totals 103.5 GW, can be broken down as follows: nuclear (11.966 GW), lignite (20.694 GW), coal (25.488 GW), gas (35.751 GW), oil (2.204 GW), biomass (7.179 GW), and waste (0.188 GW) [source: European Network of Transmission System Operators (ENTSO-E) transparency platform]. The peak load of the system in 2014 was 73.218 GW [source: European Network of Transmission System Operators (ENTSO-E) transparency platform]. The pumped hydro storage and pumping capacity amounts to 5.867 GW, while the energy storage capacity of pumped hydro units in Germany amounts to 38 GWh [source: European Energy Exchange (EEX) transparency platform].

The technical specifications of 292 conventional units are available from a commercial database. Reserve requirements for German primary, secondary and tertiary reserve are obtained from regelleistung.net. Load and import data are fixed to their

TABLE I
PERFORMANCE COMPARISON AMONG DIFFERENT CONFIGURATIONS
OF SDDP SETTINGS

Nodes K	Samples M	Iter. I	Run time (hrs)	Gap (%)	Cost (10^3€)
4	50	10	2.3	1.8	18583
10	50	10	4.3	2.8	18545
20	50	10	10.5	3.1	18565
100	5	10	4.6	3.2	18733

The performances is evaluated using 100 samples on a 100-node lattice.

historical values, which are obtained from the ENTSO-E transparency platform. Pumped hydro resources have a round-trip efficiency of 76.5%. It is assumed that fast response units are available at every bus. The marginal cost of these units is assumed to be piecewise constant, and ranges from 100 €/MWh up to 500 €/MWh. The model consists of 228 buses and 312 lines.

A. System Settings

The case study is conducted in two steps: (i) A unit commitment problem is solved with a weekly horizon. The weekly horizon prevents boundary effects in the unit commitment decision, and is justified by the fact that the German system conducts a weekly reserve capacity auction. The unit commitment model is implemented with an hourly time step, and is populated with data for the German system from September 22, 2014 until September 28, 2014. The model accounts for primary, secondary and tertiary reserve requirements given a deterministic forecast of renewable energy production. This centralized unit commitment model approximates the sequential clearing of reserve capacity followed by the day-ahead exchange of power in the Central Western European day-ahead power exchange [31]. Interactions of Germany with neighboring zones [32] are not accounted for in this paper, in order to focus on the management of real-time renewable supply uncertainty. (ii) Once unit commitment decisions are fixed, the economic dispatch model is solved for the middle of the week (Thursday) with a horizon of 24 hours and a time step of 15 minutes.

The lattice of the renewable supply stochastic process is discretized in 10 levels. Thus, the multi-stage stochastic economic dispatch lattice consists of 96 stages and 10 nodes per stage. The renewable production model is calibrated using renewable supply data sourced from the ENTSO-E transparency platform for 2013 and 2014. Renewable production is distributed in each region according to its average historical profile. The sensitivity of the solution to spatial correlations of renewable supply is discussed in Section IV-D.

The choice of using 10 nodes per stage strikes an acceptable balance between run time, convergence of the solution, and accuracy of the stochastic model. Table I presents the run time and performance of various configurations of SDDP in terms of lattice size and the number of samples used in the forward pass, for a fixed number of forward-backward iterations of the algorithm. The performance of each solution is evaluated using




Fig. 3. The evolution of lower and upper bounds of the SDDP algorithm.

100 samples on a 100-node lattice. One observes that the 4-node instance performs well in terms of optimality gap, however the representation of uncertainty is coarse-grained, and results in inferior cost performance. The 20-node and 100-node instances better represent uncertainty, however they do not achieve as good an optimality gap as the 10-node instance. Moreover, the run time of the 20-node instance with 50 forward samples is increased significantly. It is worth noting that one encounters a wide range of SDDP settings in the literature [4], [27], [28], [33] (node counts K ranging from 2 to 100, iterations I ranging from 5 to 3000, forward samples M ranging from 1 to 50, horizon steps H ranging from 5 to 120), and that the choice of these setting interacts with the convergence of the algorithm.

Three policies are compared: (i) a wait-and-see policy which manages pumped hydro resources by perfectly anticipating the evolution of renewable supply over the day; (ii) the stochastic programming policy obtained by SDDP; and (iii) a deterministic equivalent dispatch, whereby the pumped hydro schedule is fixed to the output of the day-ahead unit commitment model, and conventional thermal units are dispatched at minimum cost in real time while respecting transmission constraints.

Fig. 3 presents the evolution of the bounds generated by the SDDP algorithm when solving for the stochastic programming policy. The lower bound is indicated by the dotted line. The solid line indicates the estimate of the upper bound, and it is enveloped by its 67% and 95% confidence intervals. Each $NLDS$ of the problem consists of 1371 variables and 3560 constraints. The reported results are based on the settings indicated in the second row of Table I.

B. Performance Comparison

The average cost results are shown in Table II over 100 samples. The standard deviation of the average cost estimator⁷ is also reported, in order to verify that the reported cost differences are statistically significant. The table also reports the results of a model whereby the marginal cost of the fast-start units is set at 3000 €/MWh, so that these units are only used as a solution of last resort. These cases are indicated in the last three rows of the table.

⁷The standard deviation is indicated by σ in Table II, and is defined in footnote 3. The superior performance of the stochastic programming policy with respect to deterministic dispatch is verified with a confidence level of 94.5%.

For the model with all constraints, the stochastic programming policy results in benefits that amount to 1.2% of the deterministic dispatch cost, and perfect foresight results in benefits that amount to 0.8% of the stochastic programming cost. The perfect foresight policy utilizes fast-start units, and, although for a higher penalty the utilization of these units is decreased, it cannot be eliminated completely due to re-dispatch in real time which is necessary for relieving transmission constraints.

Fig. 4 presents the 10th, 50th and 90th percentiles of the reservoir levels for the stochastic programming policy, as well as the fixed reservoir schedule of the deterministic policy. The wide spread of the 10th and 90th percentile implies that the pumped hydro schedule can vary significantly in the stochastic programming policy. The deterministic schedule closely tracks the 50th percentile of the stochastic programming policy. The deterministic policy will pump mostly in the night hours and to a lesser extent during the peak of solar production around noon, and will produce in the morning and evening load peaks. Fig. 5 presents the dispatch of pumped hydro resources under the stochastic programming policy for the case of over- and under-forecasting renewable power supply. In the case where realized production exceeds the forecast during the morning load peak, the stochastic policy produces more aggressively, in order to create space in the reservoir for the peak solar production, which exceeds the day-ahead forecast. By contrast, in the case where the forecast exceeds the realized production, the stochastic programming policy serves less of the morning peak load than the deterministic policy through pumped hydro resources, and pumps more aggressively during the night. This is explained by the anticipation of an unfavorable renewable production outcome in the afternoon, which implies that the water in the reservoir is better used for serving the evening load peak.

By comparing the base case results of Table II to the case where the marginal cost of emergency resources is set at a very high value (3000 €/MWh), one notes that the results are sensitive to the cost of activation of fast-start resources for all policies. It is interesting to note that all policies are forced to resort to fast-start resources in real time to some extent, due to binding transmission constraints.

An optimal dispatch policy will tend to carry water in the reservoirs up to the level where its marginal value equals the expected future cost of production. This will tend to align marginal prices over all hours of the day, all outcomes and all locations. This is demonstrated in the box plots of Fig. 6, which present the distribution of LMPs for both the stochastic policy as well as the deterministic dispatch.

C. Effect of Transmission and Ramping

This section analyzes the importance of transmission constraints and ramps in the performance of each policy. In order to isolate the effect of transmission and ramping, the models without transmission constraints are re-run first (in order to capture the effect of transmission constraints) and then the models without transmission constraints and without ramp constraints are re-run (in order to isolate the incremental effect of ramp constraints). The results are presented in Table III.

TABLE II
PERFORMANCE OF PERFECT FORESIGHT, STOCHASTIC PROGRAMMING, AND THE DETERMINISTIC POLICY

	Slow unit cost (10^3 €)	Fast-start cost (10^3 €)	Total cost (10^3 €)	σ total cost (10^3 €)	Fast-start energy (MWh)	Excess energy (MWh)
Perfect Foresight	17380	850	18231	291	7204	1809
Stochastic Programming	17423	955	18378	305	7511	1855
Deterministic	17373	1221	18594	350	8688	1879
Perfect Foresight (3000)	17914	9637	27551	2626	3212	3823
Stochastic Programming (3000)	18201	10489	28690	2674	3496	5406
Deterministic (3000)	17924	13652	31575	3048	4551	4498

All values are per day. The cases indicated by '(3000)' correspond to setting the marginal cost of fast-start units to 3000 €/MWh.




Fig. 4. The percentile of the reservoir level for the stochastic programming policy, and the reservoir schedule of the deterministic policy.




Fig. 5. The dispatch of pumped hydro under the stochastic programming policy in the case where (i) realized production exceeds the forecast (top 2 panels) and (ii) the forecast exceeds realized production (lower 2 panels).

By comparing the results of Tables II and III, it can be observed that transmission constraints have a major impact on the results. In the case without transmission constraints, the benefits of stochastic programming over deterministic dispatch amount to 0.3% of the cost of deterministic dispatch, while the




Fig. 6. The locational marginal prices of the stochastic and deterministic dispatch policies. The boxes include the observations within the 25% and 75% quantiles, the crosses correspond to outliers.

TABLE III
DAILY COST PERFORMANCE FOR THE CASE WHERE TRANSMISSION AND RAMPING CONSTRAINTS ARE DROPPED

	Total cost (10^3 €)	St. dev. (10^3 €)
Foresight (No Transmission)	14797	97
Stochastic (No Transmission)	14828	100
Deterministic (No Transmission)	14867	105
Lookahead 1-step (No Transmission)	14865	105
Stochastic Hydro (No Transmission)	14830	101
Foresight (No Ramp/Transmission)	14796	97
Stochastic (No Ramp/Transmission)	14828	100
Deterministic (No Ramp/Transmission)	14856	105

benefits of perfect foresight over stochastic programming amount to 0.2% of the cost of stochastic programming.

The motivation for analyzing the effect of ramping separately is the recent introduction of flexible ramp products⁸ in various US markets, including the Midwest ISO [34] and the California ISO [35]. The results of Table III suggest that the incremental cost impact of ramp constraints is negligible. In addition, a one-step deterministic lookahead dispatch is analyzed, which is inspired by multi-period look-ahead in US real-time markets that are intended to better accommodate ramp capacity scarcity

⁸The goal of flexible ramp products is to ensure that the system is optimally pre-positioned in anticipation of multi-period and potentially uncertain ramps, without undermining the provision of other ancillary service products, and while ensuring that flexible capacity is priced properly.

[34]. This policy is referred to as ‘Lookahead 1-step’ in Table III. Compared to the deterministic dispatch, this policy exhibits a minor benefit. In fact, what is more crucial than dispatching conventional thermal units with a one-step lookahead is to dispatch pumped hydro resources by anticipating the multi-period (beyond one step) evolution of uncertainty over the day. This is shown in Table III by calculating the cost of ‘Stochastic Hydro’: a policy that uses the pumped hydro schedule of the stochastic programming solution, and dispatches conventional resources without a one-step look ahead. As shown in the table, this policy achieves almost the same cost as the stochastic programming solution. This implies that dispatching pumped hydro resources while being mindful of the multi-period evolution of the system is more crucial than dispatching conventional resources with a short look-ahead in order to prevent shortage in ramp capacity. This observation was not obvious at the outset of the study.

D. Spatial Correlations

Renewable energy production in Germany exhibits strong spatial correlations [36]. Statistical analysis reveals strong positive correlation with a certain lag between all pairs of the four regions operated by German TSOs. The lagged correlations indicate weather fronts that are moving from one part of the country to the other, and resulting in similar renewable production patterns with a difference of a few hours.

Two common approaches for accounting for spatial correlation in the SDDP hydrothermal planning literature are (i) to assume serial independence, and use observations of historical realizations of inflows for populating nodes of equal likelihood in the SDDP lattice [37], [38], and (ii) to fit an autoregressive model for each region based on historical data, estimate the pairwise correlation of the residual noise of each region through historical data, and estimate the Cholesky factorization of the resulting correlation matrix [33]. The resulting Cholesky factor is then applied to a random sample of multi-dimensional white noise, where each sample of the error corresponds to a node of equal likelihood in the SDDP lattice.

Each of these methods has evident shortcomings. The first approach tends to over-fit based on historical data. The second method attempts to estimate a very large number of correlation parameters based on very few data points, and is then using an extremely sparse support to estimate a high-dimensional process. Therefore, in attempting to model spatial correlations the modeler needs to ‘pick his/her poison’. In order to test for the sensitivity of the solution to the modeling of spatial correlations, the first alternative method for modeling uncertainty mentioned in the previous paragraph has been implemented in this paper. In particular, historical data from January 2013 until December 2014 has been used in order to identify the 10 days in the data set which were closest to the day-ahead forecast of the date of the case study (September 25, 2014). The data has then been used in order to draw samples of the forecast error for each of the four regions operated by the different German TSOs (data for solar and wind forecast errors are only available by TSO operating region), and the forecast error has been spread evenly among the buses of each region. The 10 samples of forecast error data have

TABLE IV
THE SIZE OF THE *NLDS* AND THE RUN TIME OF THE SDDP ALGORITHM FOR THE MODELS REPORTED IN SECTION IV

	<i>NLDS</i> variables	<i>NLDS</i> constraints	Run time (hr)
Full	1371	3560	4.3
No Transmission	832	2272	3.3
No Ramp/Transmission	832	1880	2.8

been used as the 10 nodes of the SDDP lattice, with each node assumed to be occurring with equal probability. Presumably, this data manipulation captures spatial correlation, although, as argued earlier, it tends to over-fit the lattice to historical observations. The resulting SDDP model is in fact computationally easier to solve, because the state space no longer includes past realizations of forecast errors. The L1-norm of the change in the value functions over time stages averages at 7.5%, and is therefore non-negligible, but remains below 10% for all time stages. The comparison of the long-run cost performance of different approaches for modeling spatial correlations would require rolling planning conditional on day-ahead forecasts, and is out of the scope of the present study.

In conclusion, approaches towards capturing spatial correlation which have been proposed in the literature can be implemented in the framework of the model proposed in this paper, and may result in non-negligible impacts on the obtained value functions. However, these methods involve inevitable modeling compromises in order to overcome the curse of dimensionality.

V. PRACTICAL IMPLEMENTATION

A. Off-Line Computation of Value Functions

The run time of the SDDP algorithm on a laptop with a 4-core i7-4720HQ processor and 8 GB of RAM using the FAST toolbox and Gurobi is shown in Table IV. The result of the SDDP algorithm is a dispatch *policy* that is contingent on the realization of uncertainty for the following day. Therefore, one can argue that the run time is acceptable since the stochastic programming policy can be solved for in the day-ahead time frame by the operator of the pumped hydro resources, and stored in memory for the following day.

The dependence of the SDDP run time on model and solver settings can be deduced as follows. Denote M as the number of Monte Carlo samples employed in the forward pass of the algorithm, H as the number of stages, K as the number of nodes per stage (assumed constant over stages), and I as the number of forward-backward iterations of the algorithm (assuming a fixed number of iterations), V as the number of variables in an *NLDS*, and C as the number of constraints. The execution of the algorithm for a single iteration requires solving $H \cdot M$ problems of the size of *NLDS* in a forward pass, and $K \cdot (H - 1) \cdot M$ such problems in a backward pass. Over I iterations, one solves a total of $(K \cdot (H - 1) \cdot M + H \cdot M) \cdot I$ *NLDS* problems. Assuming that the running time of the backward passes dominates that of the forward passes, which is reasonable for large M , the overall run time of the algorithm is

approximately proportional to the number of nodes per stage K , the number of iterations I , the number of Monte Carlo samples M , and the number of stages H .

In case pumped hydro resources are scheduled on a weekly or even monthly basis, the approach presented in this paper remains useful. A pumped hydro scheduling problem with a horizon of a few weeks or months and coarse-grained time steps can be used for deriving pumped hydro storage value functions. These derived value functions can then be used as input for the real-time multi-stage dispatch problem presented in this paper.

B. On-Line Decision Making

Real-time operations are governed by tight run time requirements. Economic dispatch decisions may be updated as frequently as every five minutes, therefore economic dispatch software should be capable of providing real-time decision support within a matter of a few seconds. SDDP can be run off-line to compute value functions, and these value functions can be used in real-time economic dispatch software for optimally deploying storage. The average run time required for solving *NLDS* and obtaining on-line dispatch decisions for the experiments presented in Section IV amounts to 0.027 seconds, which is clearly acceptable from an operational perspective.

The value functions obtained from the SDDP algorithm can be used for real-time decision making even though the realized uncertainty does not correspond to *any* realization of the stochastic process in the stochastic programming formulation. For example, the deployment of pumped hydro for the *realized* rainfall on the day of the case study (September 25, 2014) was computed in 3.4 seconds (i.e., 0.035 seconds per stage). This should be contrasted to typical scenario tree formulations, where the solution of the multi-stage stochastic program is a not a value function, but instead a decision which is contingent on the state of the world. In such cases, it is unclear how the decision maker should proceed if the realized state of the world does not correspond to any of the paths of the scenario tree.

C. Integration With Market Operations

The value of the stochastic solution relative to the deterministic alternative considered in this study suggests that the adaptive management of pumped-hydro resources in intra-day and real-time markets can deliver substantial benefits. This raises an interesting issue related to incentives: who should decide about the real-time dispatch of storage resources?

In a risk-neutral setting, a multi-stage stochastic competitive equilibrium coincides with the optimal solution of a coordinated multi-stage stochastic scheduling problem [39]. In theory, therefore, risk-neutral agents who correctly anticipate the evolution of uncertainty can adjust the valuation of stored energy so as to reproduce the dispatch of a benevolent centralized planner. In practice, this notion stumbles upon the exercise of market power, the fact that agents are typically not risk neutral, and the fact that agents may not share identical views about the future.

From a market power mitigation point of view, the major challenge of managing pumped storage resources in real time is the fact that these resources are not characterized by an

intrinsic marginal cost, but instead by an opportunity cost of stored energy which depends on beliefs about the future evolution of real-time prices. Therefore, a decentralized operation of storage would create challenges in terms of market monitoring since it would not be possible for the regulator to distinguish between true scarcity and the exercise of market power⁹. This concern is especially relevant in systems with ramping scarcity caused by the integration of renewable resources [41], [42]. One extreme (the one investigated in the present study which, according to data from European markets [31], is not too far from reality for certain systems) would be to simply fix pumped hydro resources to their day-ahead schedule. At the same time, the present study demonstrates that substantial benefits can be achieved from dispatching pumped hydro resources adaptively. On another extreme, therefore, one could envision the system operator dispatching pumped hydro resources centrally, since the operator would be naturally positioned to collect all the information necessary for dispatching pumped hydro resources optimally. The results of the ‘Stochastic Hydro’ policy in Table III suggest that this solution would already deliver the majority of the potential benefits of multi-period stochastic economic dispatch.

A less centralized approach whereby market agents bid their opportunity cost for pumped hydro resources in real-time markets can still benefit from the SDDP approach proposed in this paper. The value functions obtained from the optimal solution of the centralized welfare maximization can be used for constructing real-time bidding functions for pumped hydro resources. This is a topic that will be explored in more detail in future research.

VI. CONCLUSIONS AND PERSPECTIVES

Whereas the mitigation of uncertainty in day-ahead unit commitment has received a significant amount of attention in the literature, the management of uncertainty in real-time operations remains a challenging application that has recently drawn the attention of the research community and practitioners. This paper proposes a stochastic programming approach for managing pumped hydro resources, and gains of up to 1.1% are demonstrated from stochastic dispatch relative to a deterministic policy that fixes the schedule of pumped hydro units against a deterministic forecast.

The SDDP algorithm was originally proposed for tackling multi-stage stochastic linear programs. The authors have further attempted to add the commitment of fast-start units to the model, however the results obtained so far indicate that the linear relaxation of the problem achieves only marginal benefits, and further research is required in this direction. Recent advances have extended the scope of the SDDP algorithm to multi-stage stochastic integer programs [43]. This creates the possibility for the consideration of unit commitment in the multi-stage stochastic optimization of real-time operations, which will be explored in future work.

⁹The impact of storage ownership structure on the exercise of market power has been analyzed by Sioshansi [40].




Fig. 7. An illustration of the lattice used in the SDDP algorithm.

APPENDIX

A. SDDP Lattice Notation

This section illustrates the notation introduced in Section II by example. Consider the lattice shown in Fig. 7. The lattice consists of $H = 3$ stages. Each node of the lattice is indexed by $k \in \Omega_t$, which is a possible realization of the random vector $\xi_t^T = (c_t^T, h_t^T, \text{vec}(T_t), \text{vec}(W_t))$. The set of possible paths over the lattice up to stage t is denoted as $\Omega_{[t]}$, and the ancestor of a certain $\omega_{[t]} \in \Omega_{[t]}$ is denoted as $A(\omega_{[t]}) \in \Omega_{[t-1]}$. Referring to the two first stages of the lattice of Fig. 7, $\Omega_1 = \{1\}$, and $\Omega_2 = \{1, 2\}$, since stage 1 corresponds to a single outcome and stage 2 corresponds to two possible outcomes. The set $\Omega_{[2]} = \{(1, 1), (1, 2)\}$ corresponds to the two possible paths from stage 1 to stage 2. Each path $\omega_{[2]} \in \Omega_{[2]}$ originates from the same ancestor, with $A((1, 1)) = A((1, 2)) = 1$.

Forward passes are illustrated in the left part of the figure, with $\hat{x}_t^{(i)}$ corresponding to the realization of a trial decision for stage t and Monte Carlo sample i . The sample path is indicated in bold, and corresponds to $(1, 2, 3)$. Backward passes are illustrated in the right part of the figure, with $\pi_{t,k}^{(i)}$ corresponding to the dual multipliers that are used for generating optimality cuts. The dashed box represents the set of nodes in the lattice which share a common value function approximation, $\tilde{V}_t$. Note that due to the assumption of serial independence, all nodes $k \in \Omega_t$ share the same value function approximation.

B. Pumped Hydro Model Notation

Sets

- T : set of time stages
- Ω_t : set of SDDP lattice nodes in stage t
- $\Omega_{[t]}$: set of paths on SDDP lattice from stage 1 to stage t
- G : set of generators
- N : set of buses
- K : set of transmission lines
- L : set of loads
- PH : set of pumped hydro resources
- GR : set of renewable resources

Decision variables

- c_g : cost of generator g
- p_g : power production of generator g
- ls_n : load shedding at bus n
- pd_g : pumping demand of pumped hydro unit g
- pp_g : pumping supply of pumped hydro unit g

- s_g : stored energy of pumped hydro unit g
- f_k : power flow on line k
- ps_n : production shedding in node n
- θ_n : angle of bus n
- y_g : renewable supply/renewable forecast ratio for renewable generator g

Parameters

- FC_g : fuel cost of generator g
- $A_{g,m}, B_{g,m}$: intercept and slope of segment m describing fuel cost function of generator g
- U_g : on/off status of generator g
- D_l : load of consumer l
- RF_g : renewable energy forecast of renewable generator g
- B_k : susceptance of line k
- TC_k : flow limit of line k
- Γ_g^d : pumping efficiency of unit g
- Γ_g^p : production efficiency of unit g
- $DMax_g$: pumping capacity of pumped hydro unit g
- $PMax_g$: production capacity of pumped hydro unit g
- $PMin_g$: technical minimum of conventional generator g
- RU_g/RD_g : 15-minute ramp-up / ramp-down rate of generator g
- MTL_g : maximum transition limit of generator g
- RF_g : renewable forecast of renewable generator g

ACKNOWLEDGMENT

The authors would like to thank Raphael Chabar (PSR) and Alexander Street (PUC-Rio) for helpful discussions.

REFERENCES

- [1] Y. Gu and L. Xie, "Stochastic look-ahead economic dispatch with variable generation resources," *IEEE Trans. Power Syst.*, vol. 32, no. 1, pp. 17–29, Jan. 2017.
- [2] D. Swider, "Compressed air energy storage in an electricity system with significant wind power generation," *IEEE Trans. Energy Convers.*, vol. 22, no. 1, pp. 95–102, Mar. 2007.
- [3] A. Tuohy and M. O'Malley, "Impact of pumped storage on power systems with increasing wind penetration," in *Proc. IEEE Power Energy Soc. Gen. Meeting*, 2009.
- [4] M. V. F. Pereira and L. M. V. G. Pinto, "Multi-stage stochastic optimization applied to energy planning," *Math. Program.*, vol. 52, pp. 359–375, 1991.
- [5] M. Khodayar, M. Shahidepour, and L. Wu, "Enhancing the dispatchability of variable wind generation by coordination with pumped-storage hydro units in stochastic power systems," *IEEE Trans. Power Syst.*, vol. 28, no. 3, pp. 2808–2818, Aug. 2013.
- [6] D. Pozo, J. Contreras, and E. Sauma, "Unit commitment with ideal and generic energy storage units," *IEEE Trans. Power Syst.*, vol. 29, no. 6, pp. 2974–2984, Nov. 2014.
- [7] J. P. Deane, E. J. McKeogh, and B. P. O. Gallachoir, "Derivation of intertemporal targets for large pumped hydro energy storage with stochastic optimization," *IEEE Trans. Power Syst.*, vol. 28, no. 3, pp. 2147–2155, Aug. 2013.
- [8] N. Li and K. W. Hedman, "Economic assessment of energy storage in systems with high levels of renewable resources," *IEEE Trans. Sustain. Energy*, vol. 6, no. 3, pp. 1103–1111, Jul. 2015.
- [9] R. Nürnberg and W. Römis, "A two-stage planning model for power scheduling in a hydro-thermal system under uncertainty," *Optim. Eng.*, vol. 3, no. 4, pp. 355–378, 2002.
- [10] R. Jiang, J. Wang, and Y. Guan, "Robust unit commitment with wind power and pumped storage hydro," *IEEE Trans. Power Syst.*, vol. 27, no. 2, pp. 800–810, May 2012.
- [11] Y. Wen, C. Guo, H. Pandzic, and D. S. Kirschen, "Enhanced security-constrained unit commitment with emerging utility-scale energy storage," *IEEE Trans. Power Syst.*, vol. 31, no. 1, pp. 652–662, Jan. 2016.

- [12] A. Lorca and X. A. Sun, "Adaptive robust optimization with dynamic uncertainty sets for multi-period economic dispatch under significant wind," *IEEE Trans. Power Syst.*, vol. 30, no. 4, pp. 1702–1713, Jul. 2015.
- [13] A. Lorca, X. A. Sun, E. Litvinov, and T. Zheng, "Multistage adaptive robust optimization for the unit commitment problem," *Oper. Res.*, vol. 64, no. 1, pp. 32–51, 2016.
- [14] B. Wang and B. F. Hobbs, "A flexible ramping product: Can it help real-time dispatch markets approach the stochastic dispatch ideal?" *Electr. Power Syst. Res.*, vol. 109, pp. 128–140, 2014.
- [15] B. Wang and B. F. Hobbs, "Real-time markets for flexiramp: A stochastic unit commitment-based analysis," *IEEE Trans. Power Syst.*, vol. 31, no. 2, pp. 846–860, Mar. 2016.
- [16] C. Safta, R. L.-Y. Chen, H. N. Najm, A. Pinar, and J.-P. Watson, "Efficient uncertainty quantification in stochastic economic dispatch," *IEEE Trans. Power Syst.*, vol. 32, no. 4, pp. 2535–2546, Jul. 2017.
- [17] D. Phan and S. Ghosh, "Two-stage stochastic optimization for optimal power flow under renewable generation uncertainty," *ACM Trans. Model. Comput. Simul.*, vol. 24, no. 1, 2014, Art. no. 2.
- [18] R. A. Jabr, "Adjustable robust OPF with renewable energy sources," *IEEE Trans. Power Syst.*, vol. 28, no. 4, pp. 4742–4751, Nov. 2013.
- [19] D. Bienstock, M. Chertkov, and S. Harnett, "Chance-constrained optimal power flow: Risk-aware network control under uncertainty," *SIAM Rev.*, vol. 56, no. 3, pp. 461–495, 2014.
- [20] M. Lubin, Y. Dvorkin, and S. Backhaus, "A robust approach to chance constrained optimal power flow with renewable generation," *IEEE Trans. Power Syst.*, vol. 31, no. 5, pp. 3840–3849, Sep. 2016.
- [21] T. Summers, J. Warrington, M. Morari, and J. Lygeros, "Stochastic optimal power flow based on conditional value at risk and distributional robustness," *Electr. Power Energy Syst.*, vol. 75, pp. 116–125, 2015.
- [22] R. A. Jabr, S. Karaki, and J. A. Korbane, "Robust multi-period OPF with storage and renewables," *IEEE Trans. Power Syst.*, vol. 30, no. 5, pp. 2790–2799, Sep. 2015.
- [23] J. R. Birge and F. Louveaux, *Introduction to Stochastic Programming* (Springer Series in Operations Research and Financial Engineering). New York, NY, USA: Springer-Verlag, 2010.
- [24] G. Infanger and D. P. Morton, "Cut sharing for multistage stochastic linear programs with interstage dependency," *Math. Program.*, vol. 75, pp. 241–256, 1996.
- [25] R. M. V. Slyke and R. Wets, "L-shaped linear programs with applications to optimal control and stochastic programming," *J. Appl. Math.*, vol. 17, no. 4, pp. 638–663, Jul. 1969.
- [26] A. Shapiro, "Analysis of stochastic dual dynamic programming method," *Eur. J. Oper. Res.*, vol. 209, pp. 63–72, 2011.
- [27] T. H. De-Mello, V. L. D. Matos, and E. C. Finardi, "Sampling strategies and stopping criteria for stochastic dual dynamic programming: A case study in long-term hydrothermal scheduling," *Energy Syst.*, vol. 2, no. 1, pp. 1–31, 2011.
- [28] A. Shapiro, W. Tekaya, J. P. da Costa, and M. P. Soares, "Risk neutral and risk averse stochastic dual dynamic programming method," *Eur. J. Oper. Res.*, vol. 224, pp. 375–391, 2013.
- [29] F. G. Cabral, "A proposal for a multiplicative autoregressive multivariate periodic model for the generation of inflow scenarios applicable to the planning model of the Brazilian electricity sector," Ph.D. dissertation, Dept. Mech. Eng., Federal Univ. Rio de Janeiro, Rio de Janeiro, Brazil, 2016.
- [30] P. Pinson, H. Nielsen, J. K. Moller, H. Madsen, and G. N. Kariniotakis, "Non-parametric probabilistic forecasts of wind power: Required properties and evaluation," *Wind Energy*, vol. 10, pp. 497–516, 2007.
- [31] A. Papavasiliou and Y. Smeers, "Remuneration of flexibility using operating reserve demand curves: A case study of Belgium," *The Energy Journal*, vol. 38, pp. 105–135, 2017.
- [32] I. Aravena and A. Papavasiliou, "Renewable energy integration in zonal markets," *IEEE Trans. Power Syst.*, vol. 32, no. 2, pp. 1334–1349, Mar. 2017.
- [33] M. V. F. Pereira, "Recent applications of multistage stochastic optimization to power system planning and operations," in *Proc. CORE@50 Conf.*, Université Catholique de Louvain, Louvain la Neuve, Belgium, 2016.
- [34] N. Navid and G. Rosenwald, "Market solutions for managing ramp flexibility with high penetration of renewable resource," *IEEE Trans. Sustain. Energy*, vol. 3, no. 4, pp. 784–790, Oct. 2012.
- [35] "Flexible ramping product, revised draft final proposal," California ISO, Folsom, CA, USA, Tech. Rep., 2015. [Online]. Available: http://www.optimization-online.org/DB_HTML/2016/05/5436.html
- [36] A. Papavasiliou, S. Oren, and I. Aravena, "Stochastic modeling of multi-area wind power production," in *Proc. 48th Hawaii Int. Conf. Syst. Sci.*, 2015, pp. 2616–2626.
- [37] A. Brigatto, A. Street, and D. M. Valladao, "Assessing the cost of time-inconsistent operation policies in hydrothermal power systems," *IEEE Trans. Power Syst.*, vol. 32, no. 6, pp. 4541–4550, Nov. 2017.
- [38] A. Street, A. Brigatto, and D. M. Valladao, "Co-optimization of energy and ancillary services for hydrothermal operation planning under a general security criterion," *IEEE Trans. Power Syst.*, vol. 32, no. 6, pp. 4914–4923, Nov. 2017.
- [39] A. Philpott, M. Ferris, and R. Wets, "Equilibrium, uncertainty, and risk in hydrothermal electricity systems," *Math. Program.*, vol. 157, pp. 483–513, 2013.
- [40] R. Sioshansi, "Welfare impacts of electricity storage and the implications of ownership structure," *Energy J.*, vol. 31, pp. 173–198, 2010.
- [41] E. Moiseeva, S. Wogrin, and M. R. Hesamzadeh, "Generation flexibility in ramp rates: Strategic behavior and lessons for electricity market design," *Eur. J. Oper. Res.*, vol. 261, no. 2, pp. 755–771, 2017.
- [42] E. Moiseeva and M. R. Hesamzadeh, "Strategic bidding of a hydropower producer under uncertainty: Modified Benders approach," *IEEE Trans. Power Syst.*, vol. 33, no. 1, pp. 861–873, Jan. 2018.
- [43] J. Zou, S. Ahmed, and X. A. Sun, "Nested decomposition of multistage stochastic integer programs with binary state variables," Georgia Inst. Technol., Atlanta, GA, USA, Tech. Rep., 2016. [Online]. Available: <https://www.caio.com/Documents/RevisedDraftFinalProposal-FlexibleRampingProduct-2015.pdf>




Anthony Papavasiliou (M'06) received the B.S. degree in electrical and computer engineering from the National Technical University of Athens, Athens, Greece, and the Ph.D. degree from the Department of Industrial Engineering and Operations Research, University of California at Berkeley, Berkeley, CA, USA. He holds the ENGIE Chair at the Université catholique de Louvain, Louvain-la-Neuve, Belgium, and is also a faculty member of the Center for Operations Research and Econometrics. He has served as a consultant and an intern at N-SIDE, Pacific Gas and Electric, Quantil, Sun Run, the United States Federal Energy Regulatory Commission, the Palo Alto Research Center, and the Energy, Economics and Environment Modeling Laboratory at the National Technical University of Athens.



Yuting Mou (S'16) received the Bachelor's and Master's degrees in electrical engineering from Jilin University, Changchun, China, and Zhejiang University, Hangzhou, China, in 2012 and 2015, respectively. He is currently working toward the Ph.D. degree in the Center for Operations Research and Econometrics, Université Catholique de Louvain, Louvain-la-Neuve, Belgium.



Léopold Cambier received the B.S. degree in engineering and the M.S. degree in mathematical engineering from the Université Catholique de Louvain, Louvain-la-Neuve, Belgium. He is currently working toward the Ph.D. degree in computational and mathematical engineering in the Institute for Computational and Mathematical Engineering at Stanford University, Stanford, CA, USA. His focus is on fast algorithms for linear systems, including sparse and low-rank-based direct methods, incomplete factorizations, and iterative algorithms.



Damien Scieur received the Diploma in Mathematical Engineering in 2015 from the Université Catholique de Louvain, Louvain-la-Neuve, Belgium. He is currently working toward the Ph.D. degree at the École Normale Supérieure d'Ulm, Paris, France. He works at Inria Paris, and is the recipient of a Marie Curie Fellowship.