A Computational Framework for Uncertainty Quantification and Stochastic Optimization in Unit Commitment with Wind Power Generation

Emil M. Constantinescu, Victor M. Zavala,
Matthew Rocklin, Sangmin Lee, and Mihai Anitescu

***Abstract*—We present a computational framework for integrating a state-of-the-art numerical weather prediction (NWP) model in stochastic unit commitment/energy dispatch formulations that account for wind power uncertainty. We first enhance the NWP model with an ensemble-based uncertainty quantification strategy implemented in a distributed-memory parallel computing architecture. We discuss computational issues arising in the implementation of the framework and validate the model using real wind-speed data obtained from a set of meteorological stations. We build a simulated power system to demonstrate the developments.**

***Index Terms*—weather forecasting, wind, unit commitment, energy dispatch, closed-loop.**

NOMENCLATURE

Numerical Weather Prediction

γ	Ensemble inflation factor
$\mathcal{M}$	Numerical weather prediction model
M_x	Number of model states
N_S	Number of ensemble members
$\mathbf{Q}$	Covariance of the model error
U, V, T	Horizontal wind components and the temperature atmospheric fields
$\mathbf{V}$, $\mathbf{C}$	Covariance and correlation matrices of the initial ensemble
x	Atmospheric field
$\overline{x}$, $\mathbf{S}^2$	Ensemble sample average and covariance matrix
x_{NARR}	Atmospheric state reconciled with observations

Unit Commitment

a_j,b_j	Coefficients of production cost function of thermal unit j
cc_j,hc_j,t_j^{cold}	Startup cost function coefficients of thermal unit j
$c_{j,k}^d$	Shutdown cost of thermal unit j in period k
$c_{j,k}^p$	Production cost of thermal unit j in period k
$c_{j,k}^u$	Startup cost of thermal unit j in period k
C_j	Shutdown cost of thermal unit j
D_k	Load demand in period k
DT_j	Minimum down time of thermal unit j
K_j^t	Cost of interval t of the stairwise startup function of thermal unit j
N	Number of thermal units
N_S	Number of wind-power scenarios
N_{wind}	Number of wind units
ND_j	Number of intervals of the stairwise startup function of thermal unit j
$\nu_{j,k}$	On/off state of thermal unit j in period k
$p_{s,j,k}$	Power output of thermal unit j in period k and scenario s
$p_{s,j,k}^{wind}$	Forecasted power output of wind unit j in period k and scenario s
$p_{s,j,k}^{wind,true}$	Observed power output of wind unit j in period k and scenario s
$\overline{P}_j$	Maximum power output of thermal unit j
$\underline{P}_j$	Minimum power output of thermal unit j
R_k	Reserve in period k
RD_j	Ramp-down limit of thermal unit j
RU_j	Ramp-up limit of thermal unit j
SD_j	Shutdown limit of thermal unit j
SU_j	Startup limit of thermal unit j
T	Number of periods
UT_j	Minimum up time of thermal unit j

Inference Analysis

$\mathbf{A}, b$	Coefficients of first-stage constraints
d	Coefficients of first-stage cost
$\hat{f}_{N_S}^j$	Suboptimal sample average cost for batch j
$L_{N_S,M}, s_{L,N_S,M}^2$	Mean and variance of lower bound
M	Number of data batches
q	Coefficients of second-stage cost
$\mathcal{Q}, Q$	Second-stage cost and realization
$\mathbf{T}, \mathbf{W}$	Coefficients of second-stage constraints
$U_{N_S,M}, s_{U,N_S,M}^2$	Mean and variable of upper bound
$\hat{v}_{N_S}^j$	Optimal sample average cost for batch j
ξ_s	Realization s of random variable
y	Second-stage decision variables for scenario k
z	First-stage decision variables

E.M. Constantinescu, V.M. Zavala and M. Anitescu are with the Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439, USA. E-mail: {emconsta, vzavala, anitescu}@mcs.anl.gov.

M. Rocklin is with the Department of Computer Science, University of Chicago, 1100 East 58th Street, Chicago, IL 60637, USA.

S. Lee is with Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA.

I. INTRODUCTION

Wind power is becoming worldwide a significant component of the power generation portfolio. In Europe, several countries already exhibit adoption levels in the range of 5-20% of the

total annual demand. In the U.S. an adoption level of 20% is expected by the year 2030 [1]. Such a large-scale adoption presents many challenges to the operation of the electrical power grid because wind power is highly intermittent and difficult to predict. In particular, unit commitment (UC) and energy dispatch (ED) operations are of great importance because of their strong economic impact (on the order of billions of dollars per year) and increasing emissions concerns.

Several UC studies analyzing the impact of increasing adoption levels of wind power have been performed recently. In [21], a security-constrained stochastic UC formulation that accounts for wind-power volatility is presented together with an efficient Benders decomposition solution technique. In [19], a detailed closed-loop stochastic UC formulation is reported. The authors analyze the impact of the frequency of recommitment on the production, startup, and shutdown costs. They find that increasing the recommitment frequency can reduce costs and increase the reliability of the system. None of these previous stochastic optimization studies present details on the wind-power forecast model and uncertainty information used to support their conclusions. In [12], [15], artificial neural network (ANN) models are used to compute forecasts and confidence intervals for the total aggregated power for a set of distributed wind generators. A problem with empirical (data-based) modeling approaches [5], [20], [22], however, is that their predictive capabilities rely strongly on the presence of persistent trends. In addition, they neglect the presence of spatio-temporal physical phenomena that can lead to time-varying correlations of the wind speeds at neighboring locations. Such approaches can thus result in inaccurate medium- and long-term forecasts and over- or underestimated uncertainty levels [14], [8], [13], which in turn affect the expected cost and robustness of the UC solution. A comparison between uncertainty quantification techniques with empirical and physical weather prediction models for ambient temperature forecasting is presented in [23].

In this work, we seek to exploit recent advances in numerical weather prediction (NWP) models to perform UC/ED studies with wind-power adoption. The use of physical models is desirable because consistent and accurate uncertainty information can be obtained [13]. In a previous study, we have found that NWP models allow one to obtain much tighter uncertainty intervals of temperature forecasts that translate into lower operating costs in building systems [23]. On the other hand, we have also found that the practical capabilities of NWP models are limited. One of the major limiting factors is their computational complexity. For instance, performing data assimilation every hour at a high spatial resolution is currently not practical. In addition, extracting uncertainty information from NWP models quickly becomes intractable from the point of view of both simulation time and memory requirements. The question is: *From an operational point of view, how suitable and practical are the forecasting capabilities of state-of-the-art NWP models?* This is an important question because NWP models are expected to be used to make real-time operational decisions with important economic implications. To analyze this issue, we present a framework that integrates the Weather Research and Forecast (WRF) model with a closed-loop stochastic UC/ED formulation. In particular, we are interested in analyzing computational issues and the effects of wind uncertainty on UC/ED operations.

Arguably, more sophisticated hybrid methods that combine both NWP wind speed forecasts and empirical models are needed to map the resolution of NWP forecasts down to a specific domain and to account for system-specific characteristics (e.g., power curves, orography) [13], [16], [6]. We point out, however, that our approach offers several advantages over previous work involving wind forecast, such as [16], [6]. The fact that we have control over both the UC/ED model and the WRF model allows us to refine the wind forecast with uncertainty as needed. In particular, as presented in Sec. IV-A, we can run the WRF at higher resolution than the data in [16], [6], and we also have control over the number of scenarios used. The latter capability can have a large impact on the UC/ED solution feasibility and efficiency and can be used in conjunction with the confidence estimation described in Sec. III-D either to increase the number of scenarios in order to improve the uncertainty precision, if needed, or to use a more calculated conservative solution. Full quantification of these benefits is an important medium-term goal of our project. The goal of this work, however, is to describe the benefits of the integration framework and for UC/ED problems.

We model the uncertainty of the wind-speed forecasts using a sampling technique that generates an ensemble of the future realizations in the targeted geographical region. The ensembles are obtained by using a scalable implementation on a distributed-memory parallel computing and are sent to a stochastic UC/ED problem. A resampling technique is developed to assess the quality of the stochastic UC/ED solutions. We validate the forecasts and spatial correlations using real wind-speed data obtained from a set of meteorological stations. We also perform an economic analysis of the impact of increasing adoption levels of wind power. The novelty of our work lies in the integration of uncertainty quantification and stochastic optimization, topics normally analyzed independently. From this integration, we can analyze the economic effects of forecast accuracy and uncertainty bounds. An addition novelty of our work is a computational analysis of WRF, which is important in order to understand its limitations and capabilities in an operational setting.

The paper is structured as follows. Section II presents details on the WRF model and on uncertainty quantification. Section III describes the stochastic unit commitment formulation and presents a resampling technique used to perform inference analysis. Section IV presents numerical validation results for WRF and the closed-loop UC simulation results. We conclude with a summary and directions for future work.

II. Wind Forecast and Uncertainty Estimation Using WRF

In this section, we describe the procedures used to forecast the wind speed using WRF. We present in detail the ensemble initialization and restarting procedures required in an operational framework. The WRF model [17] is a state-of-the-art numerical weather prediction system designed to serve both operational forecasting and atmospheric research needs. WRF is the result of a multi agency and university effort to build a highly parallelizable code that can run across scales ranging from large-eddy to global simulations. WRF has a comprehensive description of the

atmospheric physics that includes cloud parameterization, land-surface models, atmosphere-ocean coupling, and broad radiation models. The terrain resolution can go up to 30 seconds of a degree (less than 1 km^2).

To initialize the NWP simulations, we use reanalyzed fields, that is, simulated atmospheric states reconciled with observations, because the entire atmospheric state space is required by the model as initial conditions whereas only a small subset of the state space is available through measurement at any given time [9]. In particular, we use the North American Regional Reanalysis (NARR) data set that covers the North American continent (160W-20W; 10N-80N) with a resolution of 10 minutes of a degree, 29 pressure levels (1000-100 hPa, excluding the surface), every three hours from 1979 until present. We use an ensemble of realizations to represent uncertainty in the initial (random) wind field and propagate it through the WRF nonlinear model. The initial ensemble is obtained by sampling from an empirical distribution, a procedure similar to the National Centers for Environmental Prediction (NCEP) method [10]. In the following sections we describe in more detail the procedures needed for generating the forecast and its uncertainty. A similar approach is presented in [23].

1) Ensemble Initialization: In a normal operational mode, the NWP system evolves a given state from an initial time t_0 to a final time t_F. The initial state is produced from past simulations and reanalysis fields. Because of observation sparseness in the atmospheric field and the incomplete numerical representation of its dynamics, the initial states are not known exactly and can only be represented statistically. Therefore, we use a distribution of the initial conditions to describe the confidence in the knowledge of the initial state of the atmosphere. We assume a normal distribution of the uncertainty field of the initial state, a typical assumption in weather forecasting. The distribution is centered on the NARR field at the initial time, the most accurate information available. In other words, the expectation is exactly the NARR solution. The second statistical moment of the distribution described by the covariance matrix $\mathbf{V}$ is approximated by the sample variance or pointwise uncertainty and its correlation, $\mathbf{C}$. The initial N_S-member ensemble field $x_s^{t_0} := x_s(t_0)$, $i \in \{1 \dots N_S\}$, is sampled from $\mathcal{N}(x_{\text{NARR}}, \mathbf{V})$:

$$x_s^{t_0} = x_{\text{NARR}} + \mathbf{V}^{\frac{1}{2}} \xi_s \,,\ \xi_s \sim \mathcal{N}(0, I) \,,\ s \in \{1 \dots N_S\} \,, \quad (1)$$

where $\mathbf{C} = \mathbf{V}_{ij} / \sqrt{\mathbf{V}_{ii}\mathbf{V}_{jj}}$ and $\mathbf{V}_{ii}$ is the variance of variable i. This is equivalent to perturbing the NARR field with $\mathcal{N}(0, \mathbf{V})$. That is, $x_s = x_{\text{NARR}} + \mathcal{N}(0, \mathbf{V})$. In what follows, we describe the procedure used to estimate the correlation matrix.

2) Estimation of the Correlation Matrix: In weather models the correlation structure typically is localized in space. Therefore, in creating the initial ensemble one needs to estimate the spatial scales associated with each variable. To obtain these spatial scales, we build correlation matrices of the forecast errors using the WRF model. These forecast errors are estimated by using the NCEP method [10], which is based on starting several simulations staggered in time in such a way that, at any time, two forecasts are available. In particular, we run a month of day-long simulations started every twelve hours so that every twelve hours we have two forecasts, one started one day before and one started half-a-day before. The differences between two staggered simulations is denoted as $d_{ij} \in \mathbb{R}^{N \times (2 \times 30\text{days})}$, that is, the difference at the ith point in space between the jth pair of forecasts, where N is the number of points in space multiplied by the number of variables of interest. We can then define ϵ_s as the ith row, each of which correspond to the deviations for a single point in space. Therefore, the covariance matrix can be approximated by $\mathbf{V} \approx \mathbf{d}\mathbf{d}^{\mathbf{T}}$. Calculating and storing the entire covariance matrix are computationally intractable. Consequently, we describe the correlation distance at each vertical level and for each variable by two parameters representing the East-West and North-South directions. This approach captures the Coriolis effect and the Earth rotation, as well as faster and larger-scale winds in the upper atmosphere. We assume that correlations and winds are roughly similar in nature across the continental U.S. This process is repeated in the vertical direction. To create the perturbations from these length scales, we take a normally distributed noisy field and apply Gaussian filters in each direction with appropriate length scales to obtain the same effect as in (1).

3) Ensemble Propagation through the WRF Model: The initial state distribution is evolved through the NWP model dynamics. The resulting trajectories can then be assembled to obtain an approximation of the forecast covariance matrix:

$$x_s^{t_F} = \mathcal{M}_{t_0 \to t_F}\left(x_s^{t_0}\right) + \eta_s(t) \,,\ s \in \{1 \dots N_S\} \,, \quad (2)$$

where $x_s^{t_0} \sim \mathcal{N}(x_{\text{NARR}}, \mathbf{V}^{t_0})$, $\eta_s \sim \mathcal{N}(0, \mathbf{Q})$, and $\mathcal{M}_{t_0 \to t_F}(\bullet)$ represents the evolution of the initial condition through the WRF model from time t_0 to time t_F. The initial condition is perturbed by the additive noise η that accounts for the various error sources during the model evolution. An analysis of the covariance propagation through the model is given in [23].

In this study, we assume that the numerical model (WRF) is perfect, that is, $\eta \equiv 0$, and given the exact real initial conditions, the model produces error-free forecasts. For long prediction windows, this is a strong assumption. In this study, however, we restrict the forecast windows to no longer than one day ahead, thus making this assumption reasonable.

4) Accounting for Error Underestimation: In an operational setting, observations become available periodically and can be assimilated in the atmospheric state. In order to account for the new information, the ensemble needs to be recentered on the new reanalyzed field. In our example, we consider 12-hour windows between restarts. However, the errors given by the ensemble variance may be over- or underestimated because of simulation and sampling errors. In other words, the ensemble statistics may diverge from the true statistics. Therefore, the error levels need to be re-estimated before each initialization. Since correlations between entries in the state vector are more robustly estimated – their values are accurate under fewer assumptions compared to variance – by our approach [9], [23], variance is the only parameter that needs to be adjusted. One approach is to consider the reanalyzed field x_{NARR} as the true state, for computing corrections purposes only, and require that this solution be on average within one standard deviation as given by the ensemble spread. This approach corresponds to finding a factor γ that inflates the ensemble spread about its expectation. Let us consider again the ensemble x_s, $s = 1, \dots, N_S$ and the reanalyzed solution x_{NARR}. Denote by $\overline{x} = \frac{1}{N_S}\sum_{s=1}^{N_S} x_s$ the sample expectation and by $\sigma_j = \sqrt{\mathbf{S}^2{}_{jj}}$, $j \in \{1 \dots M_x\}$, the

standard deviation, where $\mathbf{S}^2 = \frac{1}{N_S-1}\sum_{s=1}^{N_S}(x_s-\overline{x})(x_s-\overline{x})^T$ is the sample covariance estimation. Then, we have

$$\gamma = \max\left(1, \min\left(\text{mean}_{U,V,T}\left(|x_{\text{NARR}} - \overline{x}|/\sigma\right), 4\right)\right),$$

where U, V, and T are the the wind-field components and the temperature, the ensemble variables under consideration. For this comparison, we consider only the first five layers, which include grid points located below 300 m. The new ensemble is then obtained by $x_s \leftarrow \overline{x} + \gamma\left(x_s - \overline{x}\right),\ s \in \{1 \dots N_S\}$. The factor is bounded between one, because the model error is underestimated in our case, and four to avoid large jumps in the solution and destabilize the NWP model. Experimentally, however, we noticed that $\gamma \approx 2$, which confirms that this approach tends to underestimate the uncertainty. This fact is not unexpected because the model error is not considered.

III. Unit Commitment and Energy Dispatch

In this section, we describe the unit commitment and energy dispatch formulation used and discuss extensions to account for wind-power uncertainty. The main idea behind our computational framework is to design a closed-loop UC/ED strategy using a stochastic programming formulation that incorporates weather forecast and uncertainty information from the WRF model and the ensemble approach described in the previous section. With this, we explore the accuracy of the forecasts and the effect of assimilating measurement information at different frequencies (WRF model reanalysis).

A. Deterministic Formulation

The UC problem has been studied extensively in literature reports. The interested reader can refer to [3], [21], [18], [7]. The UC formulation considered here is based on the mixed-integer linear programming (MILP) formulation of Carrion and Arroyo [3]. The formulation is shown below. The sets $\mathcal{T} := \{1 \dots T\}$, $\mathcal{N} := \{1 \dots N\}$, and $\mathcal{N}_{wind} := \{1 \dots N_{wind}\}$ represent the time periods, thermal units, and wind generators, respectively. The demand at each time period k is denoted by D_k, and the reserve requirement is R_k. The power output of unit j at time k is given by the continuous variable $p_{j,k}$. The expected value of the output of the wind unit j at time k, $\mathbb{E}[p^{wind}_{j,k}]$ is approximated by $\frac{1}{N_S}\sum_{s=1}^{N_S} p^{wind}_{s,j,k}$. The continuous variable $\overline{p}_{j,k}$ represents the maximum power output of unit j at time k. This variable is introduced in order to model the spinning reserves given by the differences $\overline{p}_{j,k} - p_{j,k}$. The units of all the power outputs are MW. The on/off status of unit j at time k is given by the binary variable $\nu_{j,k}$.

$$\min_{p_{j,k},\overline{p}_{j,k}\nu_{j,k}} \sum_{j\in\mathcal{N}}\sum_{k\in\mathcal{T}} c^p_{j,k} + c^u_{j,k} + c^d_{j,k} \tag{3a}$$

$$\text{s.t.} \sum_{j\in\mathcal{N}} p_{j,k} + \sum_{j\in\mathcal{N}_{wind}} \mathbb{E}[p^{wind}_{j,k}] = D_k,\ k \in \mathcal{T} \tag{3b}$$

$$\sum_{j\in\mathcal{N}} \overline{p}_{j,k} + \sum_{j\in\mathcal{N}_{wind}} \mathbb{E}[p^{wind}_{j,k}] \geq D_k + R_k,\ k \in \mathcal{T} \tag{3c}$$

$$(4) - (11).$$

The production cost for each thermal unit is approximated by using the linear model [2]

$$c^p_{j,k} = a_j\nu_{j,k} + b_j p_{j,k},\ j \in \mathcal{N},\ k \in \mathcal{T}, \tag{4}$$

where a_j and b_j are cost coefficients. To model the startup cost $c^u_{j,k}$ we use a staircase cost $K^t_j,\ t = 1, ..., ND_j$, where ND_j is the number of intervals. This leads to the following set of inequality constraints:

$$\begin{aligned} c^u_{j,k} &\geq K^t_j\left(\nu_{j,k} - \sum_{n=1}^{t}\nu_{j,k-n}\right), \\ &\qquad j \in \mathcal{N},\ k \in \mathcal{T},\ t = 1, ..., ND_j, \end{aligned} \tag{5a}$$

$$c^u_{j,k} \geq 0,\ j \in \mathcal{N},\ k \in \mathcal{T}. \tag{5b}$$

The formulation of the shutdown cost is given by

$$c^d_{j,k} \geq C_j\left(\nu_{j,k-1} - \nu_{j,k}\right),\ j \in \mathcal{N},\ k \in \mathcal{T}, \tag{6a}$$

$$c^d_{j,k} \geq 0,\ j \in \mathcal{N},\ k \in \mathcal{T}, \tag{6b}$$

where C_j is the shutdown cost of unit j. The power output of each unit at each period must satisfy the bounds

$$\underline{P}_j\nu_{j,k} \leq p_{j,k} \leq \overline{p}_{j,k},\ j \in \mathcal{N},\ k \in \mathcal{T}, \tag{7a}$$

$$0 \leq \overline{p}_{j,k} \leq \overline{P}_j\nu_{j,k},\ j \in \mathcal{N},\ k \in \mathcal{T}, \tag{7b}$$

where $\overline{P}_j$ and $\underline{P}_j$ are the maximum and minimum capacities of unit j, respectively. The thermal power outputs must also satisfy the ramp-up limits

$$\begin{aligned} \overline{p}_{j,k} &\leq p_{j,k-1} + RU_j\nu_{j,k-1} + SU_j\left(\nu_{j,k} - \nu_{j,k-1}\right) \\ &\qquad + \overline{P}_j(1 - \nu_{j,k}),\ j \in \mathcal{N},\ k \in \mathcal{T}, \end{aligned} \tag{8}$$

and the shutdown and ramp-down limits are

$$\begin{aligned} \overline{p}_{j,k-1} &\leq p_{j,k} + RD_j\nu_{j,k} + SD_j\left(\nu_{j,k-1} - \nu_{j,k}\right) \\ &\qquad + \overline{P}_j(1 - \nu_{j,k-1}),\ j \in \mathcal{N},\ k \in \mathcal{T}. \end{aligned} \tag{9}$$

Here, RD_j, RU_j, SD_j, and SU_j are the ramp-down, ramp-up, shutdown, and startup limits of unit j, respectively. The minimum uptime constraints are

$$\sum_{k=1}^{G_j}(1 - \nu_{j,k}) = 0,\ j \in \mathcal{N} \tag{10a}$$

$$\begin{aligned} &\sum_{n=k}^{k+UT_j-1} \nu_{j,n} \geq UT_j\left(\nu_{j,k} - \nu_{j,k-1}\right),\ j \in \mathcal{N} \\ &k = G_j + 1, \dots, T - UT_j + 1 \end{aligned} \tag{10b}$$

$$\begin{aligned} &\sum_{n=k}^{T}\left(\nu_{j,n} - (\nu_{j,k} - \nu_{j,k-1})\right) \geq 0,\ j \in \mathcal{N}, \\ &k = T - UT_j + 2, \dots, T, \end{aligned} \tag{10c}$$

where UT_j are the minimum uptime limits and $G_j = \min(T, (UT_j - U^0_j)\nu_{j,0})$ is the number of periods that unit j must be initially ON. The initial state of unit j is denoted by $\nu_{j,0}$ and is a fixed parameter. The minimum downtime constraints

are formulated as

$$\sum_{k=1}^{L_j} \nu_{j,k} = 0,\ j \in \mathcal{N} \tag{11a}$$

$$\sum_{n=k}^{k+DT_j-1} (1-\nu_{j,n}) \geq DT_j\big(\nu_{j,k-1}-\nu_{j,k}\big),\ j \in \mathcal{N},$$
$$k = L_j+1,\ldots,T-DT_j+1 \tag{11b}$$

$$\sum_{n=k}^{T} \big(1-\nu_{j,n}-(\nu_{j,k-1}-\nu_{j,k})\big) \geq 0,\ j \in \mathcal{N},$$
$$k = T-DT_j+2,\ldots,T\,, \tag{11c}$$

where DT_j denote the minimum downtime limits and $L_j = \min(T,(DT_j - S_j^0)(1-\nu_{j,0}))$ is the number of periods that unit j must be initially OFF. Note that it is possible to use this model to simulate the performance of the energy dispatch problem by fixing the commitment variables $\nu_{j,k}$.

B. Stochastic Programming Formulation

We extend the previous deterministic formulation by considering corrective actions on the power outputs of the thermal generators to account for the uncertainty in the wind power outputs. The problem can be cast as a two-stage stochastic programming problem similar to the ones proposed in [2], [21], [19]. The first-stage decision variables are the current thermal power outputs $p_{j,1}$, $\overline{p}_{j,1}$ and the commitment profiles over the entire planning horizon $\nu_{j,k}$. The power outputs are nonanticipatory (here and now) because it is assumed that the current wind-power outputs $p_{j,1}^{wind}$ are known and given by $p_{j,1}^{wind,true}$. To formulate the second stage, we consider multiple realizations of the wind power outputs $p_{s,j,k}^{wind}$, and we define scenario-dependent thermal power outputs $p_{s,j,k}$ and $\overline{p}_{s,j,k}$ with $k > 1$ (wait and see). Note that we do not define second-stage scenario-dependent commitment variables because we wish to keep the problem computationally tractable. The formulation of the stochastic optimization problem is

$$\min_{p_{s,j,k},\overline{p}_{s,j,k},\nu_{j,k}} \frac{1}{N_S}\sum_{s\in\mathcal{S}}\left(\sum_{j\in\mathcal{N}}\sum_{k\in\mathcal{T}} c_{s,j,k}^{p} + c_{j,k}^{u} + c_{j,k}^{d}\right) \tag{12a}$$

$$\text{s.t.} \sum_{j\in\mathcal{N}} p_{s,j,k} + \sum_{j\in\mathcal{N}_{wind}} p_{s,j,k}^{wind} = D_k,\ s\in\mathcal{S}, k\in\mathcal{T} \tag{12b}$$

$$\sum_{j\in\mathcal{N}} \overline{p}_{s,j,k} + \sum_{j\in\mathcal{N}_{wind}} p_{s,j,k}^{wind} \geq D_k + R_k,\ s\in\mathcal{S}, k\in\mathcal{T} \tag{12c}$$

$$p_{s,j,1} = p_{1,j,1}\ s\in\mathcal{S},\ j\in\mathcal{N} \tag{12d}$$

$$\overline{p}_{s,j,1} = \overline{p}_{1,j,1}\ s\in\mathcal{S},\ j\in\mathcal{N} \tag{12e}$$

$$(4)-(9),\quad s\in\mathcal{S}$$

$$(10)-(11),$$

where $\mathcal{S} := \{1\ldots N_S\}$. The ramp and power-limit constraints are defined over each scenario, $s \in \mathcal{S}$ where $p_{j,k} \leftarrow p_{s,j,k}$ and $\overline{p}_{j,k} \leftarrow \overline{p}_{s,j,k}$. The nonanticipativity constraints for the power outputs in the first time step are given by equations (12d) and (12e). For the known wind power outputs we set $p_{s,j,1}^{wind} \leftarrow p_{j,1}^{wind,true}$. Note that if the stochastic formulation is able to capture the uncertainty of the wind power accurately, the reserve requirements can be reduced to less conservative levels or even be removed. Note also that one can solve a closed-loop *stochastic* dispatch problem by fixing the commitment actions.

C. Closed-Loop Implementation

To simulate the closed-loop performance of the power system, we consider a *rolling-shrinking* horizon approach. The starting rolling time is reset to one each time new fossil fuel and electricity price information is obtained from the commodity and day-ahead markets. The latter information is available from the independent system operators. This period is assumed to be $T = 24$ hours. At the start of the rolling time, we assume that the wind-power forecast becomes available from WRF for the next 24 hours. At this point, the stochastic unit commitment problem is solved by using the current wind-power outputs $p_{s,j,\ell}^{wind} \leftarrow p_{j,\ell}^{wind,true}$ and the future forecasts $p_{s,j,k}^{wind},\ k = \ell+1,...,T$, where ℓ is the current time step. The solution of this problem gives the commitment profiles $\nu_{j,k}$ over the 24-hour rolling horizon. At each step inside the rolling horizon $\ell = 2,...,T$, the horizon is shrunk by one time step $T \leftarrow T-\ell$, and the stochastic energy dispatch is solved over the remaining horizon with the new, true wind power but the same WRF forecasts of the current day. Each of these shrinking horizon problems gives the current power outputs $p_{j,\ell}$ and $\overline{p}_{j,\ell}$ at current time ℓ.

D. Inference Analysis

In the above stochastic formulation, the wind-power outputs are assumed to have a probability distribution $\mathbb{P}$. In most stochastic optimization studies this distribution is assumed to be known. As seen in Section II, obtaining this distribution is part of the modeling task. Since many different forecast models (autoregressive, ANN, physics based) can be used to construct the error distribution, there is not a unique distribution. From a practical point of view, we expect that such a distribution is able to encapsulate the actual realizations of wind power and has tight confidence intervals. We model the wind speed distribution by propagating an assumed Gaussian distribution of the initial state conditions through the WRF model. Because of the complexity of the model, we are limited to a single batch of a few (less than a hundred) samples. From a stochastic optimization point of view, this is an issue because we are not solving the problem with the full distribution. Consequently, we must perform an inference analysis to assess the quality of the solution.

1) Computation of Confidence Intervals: The two-stage stochastic UC problem with fixed binary variables can be expressed in the following abstract form [11]:

$$\min_{z\geq 0}\ d^T z + \mathcal{Q}(z),\ \text{s.t. } \mathbf{A}z = b. \tag{13}$$

Here, z are the first-stage decision variables, and $\mathcal{Q}(z) = \mathbb{E}\left[Q(z,\xi)\right]$ is the second-stage cost. We assume that the probability distribution $\mathbb{P}$ of ξ has finite support; that is, ξ has a finite number of scenarios $\{\xi_1,...,\xi_K\}$ with probabilities $\pi_s \in (0,1)$, so we have $\mathcal{Q}(z) = \frac{1}{K}\sum_{s=1}^{K} Q(z,\xi_s)$, where

$$Q(z,\xi_s) = \min_{y_s\geq 0}\ q^T y_s,\ \text{s.t.}\ \mathbf{T}z + \mathbf{W}y_s = \xi_s,\ s = 1,...,K. \tag{14}$$

Here, y_s are the second-stage decision variables, and ξ_s are the realizations of the wind-power outputs. Since K is a very large number, it is impractical to solve the stochastic problem exactly. Therefore, given a fixed number of realizations $N_S \ll K$ from WRF, we solve the approximate problem,

$$\min_{z\geq 0} \; d^T z + \frac{1}{N_S}\sum_{s=1}^{N_S} Q(z,\xi_s), \text{ s.t. } \mathbf{A}z = b. \tag{15}$$

This smaller problem is known as a sample-average-approximation (SAA) of the original problem (13), which is usually computationally intractable. We seek to estimate lower and upper bounds of the true optimal solution v^* (using the entire set of K realizations) and their corresponding confidence intervals. Here we use the methodology put forth in [11]. A lower bound can be estimated generating $j = 1, ..., M$ batches, each of size N_S, and we can then solve (15) for each batch. If we denote as $\hat{v}^j_{N_S}$ the optimal cost of each SAA problem, we can estimate the lower bound as $L_{N_S,M} = \frac{1}{M}\sum_{j=1}^{M} \hat{v}^j_{N_S}$. The sample variance estimator is given by $s^2_{L,N_S,M} = \frac{1}{M-1}\sum_{j=1}^{M}\left(\hat{v}^j_{N_S} - L_{N_S,M}\right)^2$. The mean and variance can be used to construct confidence intervals of the lower bound. To estimate the upper bound, we pick a given value for the first-stage variables $\hat{z}$ and generate a new set of $j = 1, ..., M$ batches of data. We then evaluate (13), leading to $\hat{f}^j_N(\hat{z})$. Note that each evaluation involves the solution of the second-stage problem (14). As before, we have the mean $U_{N_S,M} = \frac{1}{M}\sum_{j=1}^{M}\hat{f}^j_{N_S}$ and variance $s^2_{U,N_S,M} = \frac{1}{M-1}\sum_{j=1}^{M}\left(\hat{f}^j_{N_S} - U_{N_S,M}\right)^2$.

2) Weighted Average Sampling: The inference analysis task requires multiple batches of realizations. As expected, obtaining these from WRF is not practical. Here, we present a heuristic resampling technique to avoid this limitation. To create new time series from the existing batch of WRF realizations, we express a new realization as a weighted average of the available ones. Suppose the WRF model is $x(t) = \mathcal{M}(t, x(0))$, where $x(t)$ is the state vector at time t. If we are given N_S samples x_j and we can write $x(0) = \sum_{j\in\mathcal{S}} w_j x_j(0)$, the propagation of $x(0)$ is $x(t) = \mathcal{M}(t, x(0)) = \mathcal{M}(t, \sum_{j\in\mathcal{S}} w_j x_j(0))$. Assuming the variance of the samples is small, we can write $x_j(0) = \bar{x}(0) + \epsilon_j(0)$. We justify the computation of weighted averages of the time series by observing that

$$\begin{aligned} x(t) &= \mathcal{M}\left(t, \bar{x}(0) + \sum_j w_j\epsilon_j(0)\right) \\ &\approx \mathcal{M}(t,\bar{x}(0)) + \sum_j w_j \frac{\partial \mathcal{M}}{\partial x}(t,\bar{x}(0))\epsilon_j(0) \\ &\approx \sum_j w_j \mathcal{M}\Big(t, \bar{x}(0) + \epsilon_j(0)\Big) = \sum_j w_j x_j(t). \end{aligned}$$

In other words, the weighted average approximates, to first order, the nonlinear propagation of weighted samples of the initial conditions. The weights are chosen to be Gaussian near the unit vectors in the standard basis on a hyperplane $\sum_{j\in\mathcal{S}} w_j = 1$ in the w space.

IV. Integrative Study

In this section we integrate the wind-speed forecasts produced by WRF by following the procedure described in Section II with the stochastic unit commitment/energy dispatch formulations described in Section III. The entire computational framework is sketched in Figure 1.

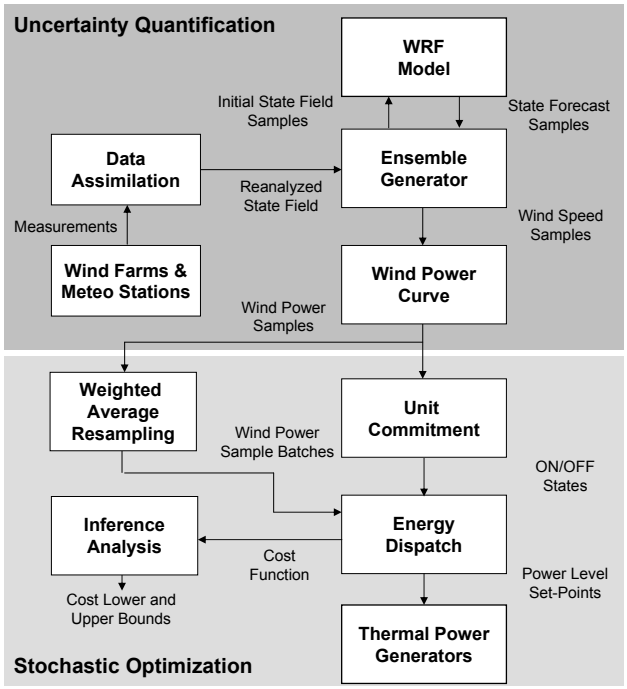

Fig. 1. Schematic representation of computational framework.

A. Wind Forecast and Uncertainty Quantification

We use the WRF model to forecast the wind speed in a specific region that covers the state of Illinois. We set up a computational nested domain structure including a high-resolution sector that covers the target area and two additional domains of larger coverage but lower resolution. The parent domains supply the boundary conditions for the nested ones, and the largest domain has prescribed boundary conditions from coarser ones. This setup is illustrated in Fig. 2. A similar setup with one coarse domain is described in [23]. We generate six ensemble data sets, each containing the predicted wind speed for Illinois corresponding to domain # 3 in Fig. 2. Each ensemble has $N_S = 30$ members. The data is sampled every 10 minutes, and each ensemble is evolved one day ahead. The starting time of the experiment t_0 corresponds to June 1, 2006, 6:00 PM CT (local time), with each data set restarted from the reanalyzed solution at time $t_0 + (k-1)\times(12\,\text{hours})$ with $k = 1, \ldots, 6$. In other words, each data set is started at the revalidation time with 12-hour increments.

1) Validation Using Wind and Temperature Data Measurements: The weather station observations were obtained from the National Climatic Data Center (NCDC), and their locations are illustrated in Fig. 3.b. In Fig. 4 we show the wind-speed ($\pm 2\sigma$) predictions and measurements for Peru and Chicago, IL (denoted by ∇ in Fig. 3.b). Each ensemble evolves for 24 hours, and new ones are started every 12 hours. We remark that the

Grid/Size
#1 - 32 km^2 130 × 60
#2 - 6 km^2 126 × 121
#3 - 2 km^2 202 × 232

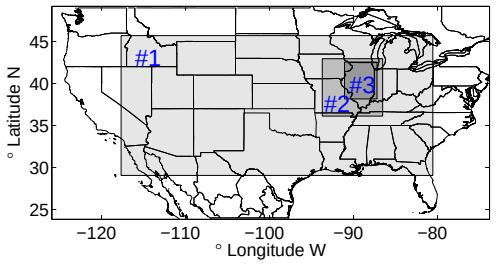

Fig. 2. Size and spatial resolution of the computational domain.

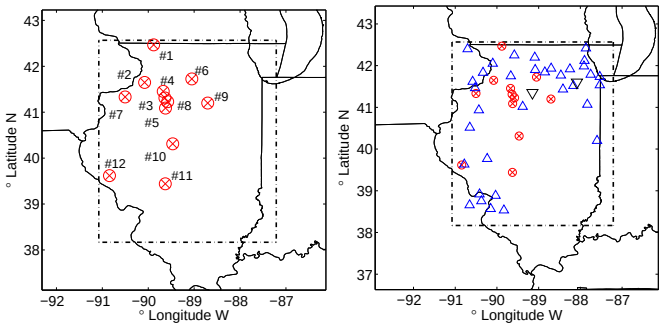

a) Wind farm location b) Weather stations

Fig. 3. Wind farms (circles) and meteorological stations (triangles) locations in Illinois.

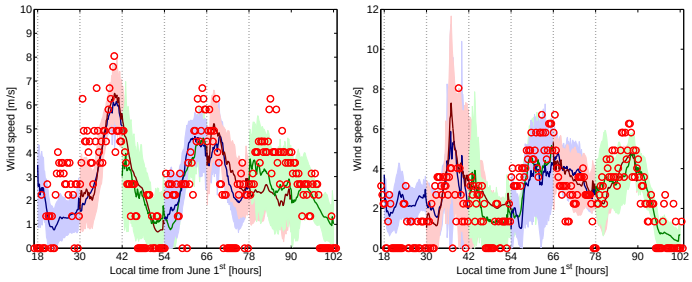

Fig. 4. Wind-speed ($\pm 2\sigma$) predictions and measurements (o) for Peru (left, RMS error=1.56, R2=0.32) and Chicago (right, RMS error=1.51, R2=0.35), IL. The vertical dashed lines denote the beginning of a new 24-hour prediction window; different colors are used to indicate ensembles started at different times.

wind-speed measurements obtained from NCDC are given in miles per hour rounded to the nearest integer. Doing so has the unfortunate effect of diminishing the wind variability and yielding more pessimistic than real validation results. Despite this, the wind-speed uncertainty intervals generated by WRF capture the trends well, with few exceptions. It is also clear from Fig. 4 that the forecasts do not improve much when updated every 12 hours instead of every 24 hours. Note that the forecasts are not improved significantly at the middle of the day, perhaps because measurements assimilated during the day are not as informative as those assimilated during the night, where the wind currents tend to be stronger. We have also observed that the wind-speed trends are much more difficult to predict than temperature trends. This point is enforced by the correlogram for the temperature and wind speed at Peru, IL shown in Fig. 6, where it is clear that the time correlations of wind speed decay more quickly than those of the temperature.

We present validation results at six active wind-farms in the state of Illinois to analyze their magnitude and correlation structure. The order of the windows goes from left to right and coincides with the wind-farm location numbering shown in Fig. 3. Currently, the power produced by wind turbines depends on the wind speed at elevations of about 40-120 meters. The wind-speed fields at these heights can be extracted from WRF. Unfortunately, the NCDC data available for validation is reported only at 10 meters. Obtaining real wind-speed data at higher altitudes requires access to proprietary databases of operational wind farms. The wind-speed fields at 10 meters above the ground for three consecutive days of June 2006 are presented in Fig. 5. The WRF realizations are able to capture the general trends of the actual observations at all locations. In addition, they are able to encapsulate the observations. Note that the wind speed is relatively low at this height. The maximum average is around 6-7 meters per second. We have found that the wind speeds reach a maximum average of around 10 meters per second at 100 meters in the studied region. In addition, we have observed that the uncertainty levels increase significantly at this height as a result of the larger range and variability. This increase is also expected because most of the wind speed data assimilated in WRF is near ground level. The 100 meters profiles are not presented here because of space restrictions. For more details, please refer to Section 4.1 in the technical report [4]. In Fig. 7 we show the spatial correlations of the wind speed for a particular wind farm on June 5, 1:50 AM, as inferred from the 30-member WRF ensemble simulation. The wind speed is highly correlated over the studied region, and it has a nontrivial spatial structure.

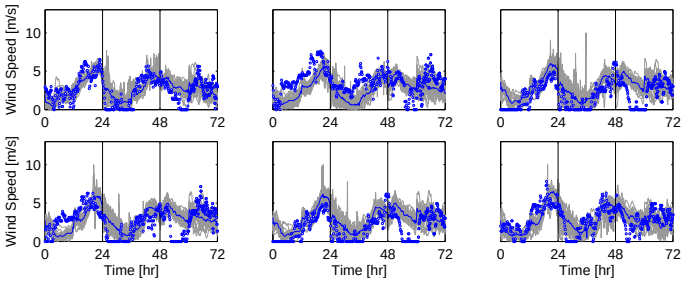

Fig. 5. Wind-speed realizations for 6 wind farm locations in Illinois at 10m and observations (dots) at nearest meteorological stations. Vertical lines represent beginning of day (12:00 AM).

2) Implementation Considerations: In this study we used version 3.1 of WRF [17]. The ensemble approach taken for estimating the uncertainty in the weather system is highly parallelizable because each scenario evolves independently through WRF. The most expensive computational element is the evolution of each sample through the WRF system. We therefore consider a two-level parallel implementation scheme. The first level is a coarse-grained task decomposition represented by each sample. A secondary finer-grained level consists in the parallelization of each sample. This approach yields a highly scalable solution. The simulations were performed on the Jazz Linux cluster at Argonne National Laboratory. Jazz (now decommissioned) had 350 compute nodes, each with a 2.4 GHz Pentium Xeon with 1.5 GB of RAM and used Myrinet 2000 and Ethernet for interconnect. Our running times given in Fig. 8 indicate that around 32 CPUs were sufficient to generate forecasts with WRF

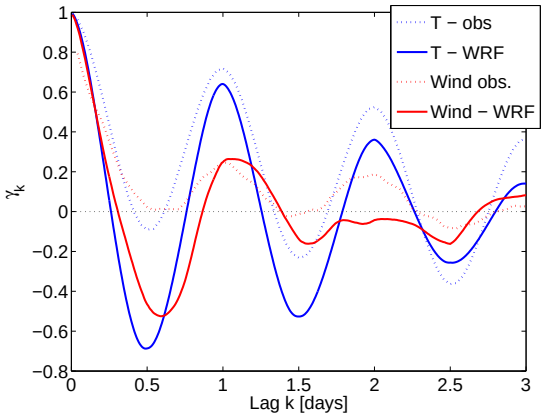

Fig. 6. Correlogram for the wind and temperature measurements and simulations at Peru, IL.

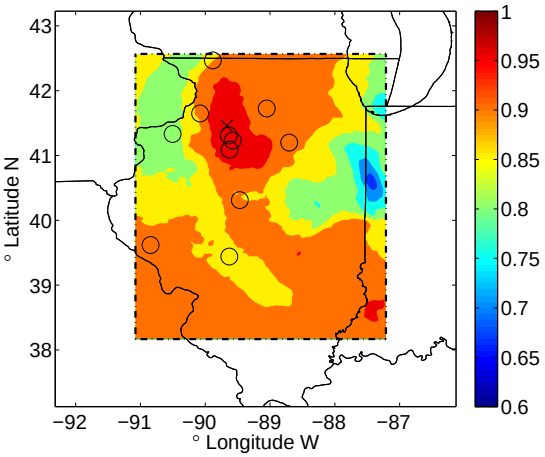

Fig. 7. Spatial correlation for the wind field for wind farm #8 on June 5, 1:50 AM, denoted by "X." The circle markers denote the other wind farms in Illinois.

in a closed-loop UC/ED setting. The times also suggest that, in order to generate forecasts every hour, one would need about 500 CPUs.

CPUs	Time [hr]
4	50
8	28
16	17
32	10

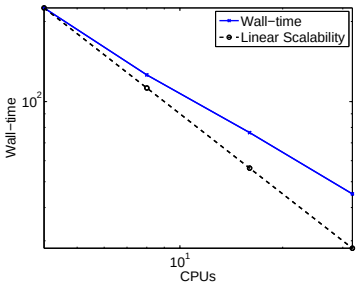

Fig. 8. Scalability of WRF on the computer cluster Jazz for 24 hours.

B. Economic Study Unit Commitment/Energy Dispatch

Because of the lack of detailed design data of thermal and wind-power units in the open literature, we have constructed an artificial simulation study. We first describe the thermal and wind-power assumptions used and then discuss our results from the simulation.

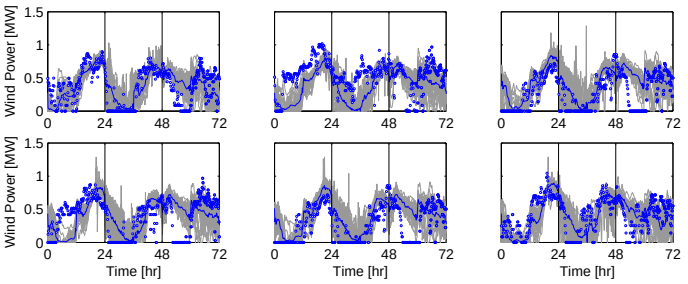

Fig. 9. Wind power realizations for 6 wind farm locations in Illinois at 10 m and observations (dots) at nearest meteorological stations. Vertical lines represent beginning of day (12:00 am).

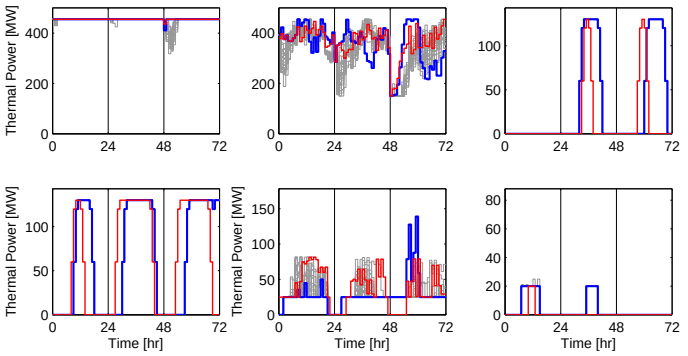

Fig. 10. Closed-loop profiles for thermal units. Solid thin line is optimal profile (with perfect information), solid thick line is stochastic UC solution, and thick gray lines are planned scenarios at the beginning of each day.

1) Power System Description: The thermal power system specifications used in this work are based on those reported in [3]. The system contains a total of 10 thermal generators with a total installed capacity of 1662 MW. The peak demand is 1326 MW. The ramp limits of the units are not reported, so we have assumed them to be 50% of the corresponding maximum capacity. The reserve requirements are assumed to be 10% of the demand. To simulate increasing level of wind power adoption, we increase the number of wind turbines at 12 existing wind farm locations in Illinois.

2) Results: To generate wind-power forecasts, we propagate the wind-speed observations and the WRF realizations at a height of *10 meters* through a typical wind-power curve with a maximum capacity of 1.5 MW. The nominal curve has a cut-in speed of 3 meters per second and reaches the rated capacity at 12 meters per second. The wind-speed observations, forecast, and ensembles used are summarized in Fig. 5. As previously mentioned, we used the height of 10 meters because the NCDC data used for validation are reported only at this level. As expected, the wind speeds are relatively low at this level, thus leading to small power outputs. Instead of using the wind speed WRF forecasts at 100 meters, we have kept the 10 meters WRF forecasts and observations and mapped these using a shifted power curve obtained by displacing the nominal cut-in speed from 3 to 2 meters per second. With this, the rated capacity is reached at around 11 meters per second. This strategy allowed us to obtain more consistent validation results for wind power compared to linear interpolation of the wind-speed observations. The resulting wind-power realizations

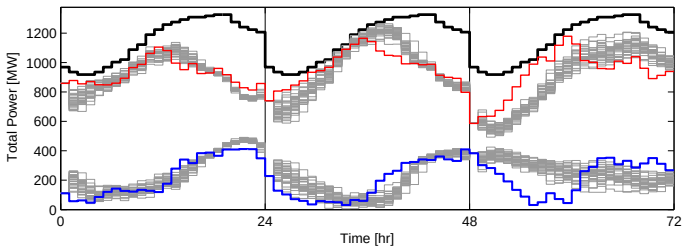

Fig. 11. Closed-loop total power profiles obtained with stochastic UC formulation. Top thick line is demand profile, medium thick line is the implemented thermal profile, gray lines are planned realizations at beginning of each day, bottom thick line is actual total wind power, and adjacent gray lines are forecast profiles.

and observations are presented in Fig. 9. The wind-power distribution is clearly affected by the nonlinear structure of the power curve, increasing the spread of the distribution. The WRF realizations are able to encapsulate the actual power observations. The largest differences are observed at the beginning of the third day.

We have run the closed-loop UC/ED system assuming a rolling horizon and a forecast frequency of 24 hours. The ED problem runs every hour. A total of 30 WRF realizations are used to solve the stochastic problem. The resulting MILP problems are implemented in AMPL and are solved with the CBC solver from the COIN-OR repository. The MILP contains 38,651 variables from which 240 are binary, 783 equality constraints, and 40,747 inequality constraints. The average solution time for the stochastic UC problem in a quad-core Intel processor running Linux is about 9 minutes in cold-start mode. The solution time of the energy dispatch problem is less than 10 seconds. The results for the 20% penetration study are presented in Figs. 10 and 11. In Fig. 10, we present the policies for the power levels of six thermal units. The solid lines represent the predicted and the realized power profiles, while the gray lines represent the forecasted realizations at the beginning of the day. We notice that the sensitivity of the power levels of some units to the uncertainty of the wind power is very small. Generators #2 and #5 are the most sensitive, while generators #3 and #4 exhibit no sensitivity. We have found that the sensitivity levels depend strongly on the design characteristics and prices of the generators. We have also found that the optimal cost of the stochastic strategy over three days of operation is only about 1% larger than that of the perfect information strategy. We also performed an inference analysis using the resampling strategy of Section III-D for the first day of operation using $M = 30$ different batches. The upper bound mean was found to be $U_{N_S,M} = \$474,064$ with variance $s^2_{U,N_S,M} = 1,082\$^2$. The lower bound mean was found to be $L_{N_S,M} = \$474,317$ with variance $s^2_{L,N_S,M} = 1,656\$^2$. Both variances are less than 0.25% of the mean cost. The smaller variances indicate that 30 WRF realizations might be sufficient to estimate the optimal cost accurately. Note, however, that the mean of the upper bound is lower than that of the lower bound. This can occur because the estimates are produced on different scenarios (the estimates would be exactly equal to each other if used with the same scenarios). It indicates the presence of a bias in the estimates which has been noted in other applications of stochastic programming [11]. This bias can be reduced using variance reduction techniques such as Latin Hypercube sampling instead of Monte Carlo sampling and by increasing the number of scenarios [11]. We have also found that updating the WRF forecast every 12 hours instead of every 24 hours does not bring important economic benefits. The reason is twofold: minor improvements in forecast accuracy, as pointed out in Section IV-A, and the properties of the power system under consideration. A different outcome could be obtained with a different generator mix.

In Fig. 11 we present the profiles of total aggregated (sum over the total units) demand, thermal power, and wind power. The solid lines represent the predicted and the realized power profiles while the gray lines represent the forecasted realizations at the beginning of the day. The top solid line is the daily demand profile, which is assumed to be constant. Note that the aggregated wind-power profile (bottom) does not follow a strong periodic trend. Nevertheless, the WRF realizations are able to encapsulate the actual profiles (solid lines) during the first two days. As a result, the optimizer is always able to satisfy the load, even for an adoption level of 20%. On the third day, however, we see a significant mismatch between the forecasted wind power and the realized one in the first 12 hours of operation. In this case, the reserves are sufficient to satisfy the load. This effect could potentially be ameliorated by inflating the initial conditions of the WRF ensembles, but it cannot be predicted *a priori*. Thus, a high frequency and adaptive inflation/resampling procedure should be added to the system. We also found that a deterministic strategy (using only the WRF forecast mean) is not able to sustain adoption levels of more than 10% even with the allocated reserves. We observed that increasing the adoption levels increases the startup and shutdown costs, but these are negligible (on the order of $10,000) with respect to the total production costs.

V. Conclusions and Future Work

We presented a computational framework for the integration of the state-of-the-art Weather Research and Forecasting (WRF) model in stochastic unit commitment/energy dispatch formulations that account for wind-power uncertainty. We extended the WRF model with a sampling technique implemented in a distributed-memory parallel computing architecture to generate uncertainty information. In addition, we developed a resampling strategy that avoids expensive WRF simulations to perform inference analysis. Our simulated commitment study indicates that using WRF forecasts and uncertainty information is critical to achieve high adoption levels with minimum reserves. Our study illustrates an operational setting with real data, pointing out several issues and limitations that are not found in idealized experiments using artificial forecasts and uncertainty information. For instance, we have not found significant benefits of updating the WRF forecasts in intra-day operations. In addition, the numerical experiments indicate that a relatively large number of CPUs are required to generate forecasts and uncertainty information at a higher frequency than 12 hours. We emphasize that the integrative framework presented here is preliminary and does not consider more detailed issues such as intra-day rescheduling of unit commitment, effects of updating wind power forecasts at higher temporal resolutions (e.g., hourly), as used in the Danish

power system. These two factors affect the value of wind power forecasts during intra-day operation. Therefore, as part of future work, we are interested in developing techniques to generate forecasts at higher spatial and temporal resolution. In addition, we are interesting in generating wind-power forecast models by fusing WRF wind-speed forecasts and operational wind-power data. Thanks to our open access to WRF, our framework is highly flexible and allows us to consider these extensions. Additionally, we are interested in dealing with networks of real size with hundreds of generators, transmission constraints, and intra-day scheduling. To solve these challenging problems, we are developing algorithms for the solution of stochastic optimization problems in parallel computing architectures. This, together, with variance reduction techniques can enable the solution of large-scale problems with small statistical errors.

ACKNOWLEDGMENTS

This work was supported by the Department of Energy, through Contract No. DE-AC02-06CH11357. We acknowledge the insightful comments of the referees, which helped us improve the quality of the paper and pointed out interesting future extensions to our work. We thank Professor Mohammad Shahidehpour for comments on an earlier version of this document. We also thank Argonne National Laboratory's LCRC for the use of Jazz cluster.

REFERENCES

[1] *20% wind energy by 2030: Increasing wind energy's contribution to U.S. electricity supply*, Tech. Report DOE/GO-102008-2567, Office of Energy Efficiency and Renewable Energy, 2009.

[2] C. C. CAROE AND R. SCHULTZ, *A two-stage stochastic program for unit commitment under uncertainty in a hydro-thermal power system*, Tech. Report SC 98-11, ZIB, 1998.

[3] M. CARRION AND J. M. ARROYO, *A computationally efficient mixed-integer linear formulation for the thermal unit commitment problem*, IEEE Transactions on Power Systems, 21 (2006), pp. 1371–1378.

[4] E. CONSTANTINESCU, V. ZAVALA, M. ROCKLIN, S. LEE, AND M. ANITESCU, *Unit commitment with wind power generation: Integrating wind forecast uncertainty and stochastic programming*, Tech. Report ANL/MCS-TM-309, Argonne National Laboratory, Mathematics and Computer Science Division Technical Memorandum, http://www.mcs.anl.gov/uploads/cels/papers/TM-309A.pdf, 2009.

[5] T. EL-FOULY, E. EL-SAADANY, AND M. SALAMA, *One day ahead prediction of wind speed using annual trends*, IEEE Power Engineering Society General Meeting, (2006), p. 7.

[6] U. FOCKEN, M. LANGE, K. MOENNICH, H. P. WALDL, H. G. BEYER, AND A. LUIG, *Short-term prediction of the aggregated power output of wind farms – a statistical analysis of the reduction of the prediction error by spatial smoothing effects*, IEEE Transactions on Power Systems, 90 (2002), pp. 231–246.

[7] J. GOEZ, J. LUEDTKE, D. RAJAN, AND J. KALAGNANAM, *Stochastic unit commitment problem*, Tech. Report RC24713, IBM, 2008.

[8] J. JUBAN, L. FUGON, AND G. KARINIOTAKIS, *Uncertainty estimation of wind power forecasts*, in Proceedings of the European Wind Energy Conference EWEC08, Brussels, Belgium, 2008.

[9] E. KALNAY, *Atmospheric Modeling, Data Assimilation and Predictability*, Cambridge University Press, 2003.

[10] E. KALNAY, M. KANAMITSU, R. KISTLER, W. COLLINS, D. DEAVEN, L. GANDIN, M. IREDELL, S. SAHA, G. WHITE, J. WOOLLEN, ET AL., *The NCEP/NCAR 40-year reanalysis project*, Bulletin of the American Meteorological Society, 77 (1996), pp. 437–471.

[11] J. LINDEROTH, A. SHAPIRO, AND S. WRIGHT, *The empirical behavior of sampling methods for stochastic programming*, Annals of Operations Research, 142 (2006), pp. 215–241.

[12] K. METHAPRAYOON, W. J. LEE, C. YINGVIVATANAPONG, AND J. LIAO, *An integration of ANN wind power estimation into UC considering the forecasting uncertainty*, IEEE Trans. Ind. Appl., 43 (2007), pp. 116–124.

[13] C. MONTEIRO, R. BESSA, V. MIRANDA, A. BOTERRUD, J. WANG, AND G. CONZELMANN, *Wind power forecasting: state-of-the-art 2009*, tech. report, INESC Porto and Argonne National Laboratory, 2009.

[14] T. PALMER, G. SHUTTS, R. HAGEDORN, F. DOBLAS-REYES, T. JUNG, AND M. LEUTBECHER, *Representing model uncertainty in weather and climate prediction*, Annual Review of Earth and Planetary Sciences, 33 (2005), pp. 163–193.

[15] V. S. PAPPALA, I. ERLICH, K. ROHRIG, AND J. DOBSCHINSKI, *A stochastic model for the optimal operation of a wind-thermal power system*, IEEE Transactions on Power Systems, 24 (2009), pp. 940–950.

[16] P. PINSON, *Estimation of Uncertainty in Wind Power Forecasting, Ph. D. thesis*, Ecole des Mines de Paris, Paris, France, 2002.

[17] W. SKAMAROCK, J. KLEMP, J. DUDHIA, D. GILL, D. BARKER, M. DUDA, X.-Y. HUANG, W. WANG, AND J. POWERS, *A description of the Advanced Research WRF version 3*, Tech. Report Tech Notes-475+STR, NCAR, 2008.

[18] S. TAKRITI, J. R. BIRGE, AND E. LONG, *A stochastic model for the unit commitment problem*, IEEE Transactions on Power Systems, 11 (1996), pp. 1497–1508.

[19] A. TUOHY, P. MEIBOM, E. DENNY, AND M. O'MALLEY, *Unit commitment for systems with significant wind penetration*, IEEE Transactions on Power Systems, 24 (2009), pp. 592–601.

[20] G. VINEEL AND A. SARKAR, *A system for real-time delivery of wind farm production forecasts*, in IEEE ISIE, 2007, pp. 1933–1937.

[21] J. WANG, M. SHAHIDEHPOUR, AND Z. LI, *Security-constrained unit commitment with volatile wind power generation*, IEEE Transactions on Power Systems, 23 (2008), pp. 1319–1327.

[22] Y.-K. WU AND J.-S. HONG, *A literature review of wind forecasting technology in the world*, (2007), pp. 504–509.

[23] V. ZAVALA, E. CONSTANTINESCU, T. KRAUSE, AND M. ANITESCU, *On-line economic optimization of energy systems using weather forecast information*, Journal of Process Control, 19 (2009), pp. 1725–1736.

Emil M. Constantinescu received his B.Sc. and M.S. degrees from the College of Automatic Controls and Computers, Bucharest Polytechnic University, Romania (2001,2002) and the Ph.D. degree from Virginia Tech (2008) in computer science. He is currently the Wilkinson Postdoctoral Fellow in the Mathematics and Computer Science Division at Argonne National Laboratory. His research interests include uncertainty quantification in weather and climate models and their applications to energy systems.

Victor M. Zavala is currently a Director's Postdoctoral Fellow in the Mathematics and Computer Science Division at Argonne National Laboratory. He received the B.Sc. degree from Universidad Iberoamericana (2003) and the Ph.D. degree from Carnegie Mellon University (2008), both in chemical engineering. He has served as a research intern at ExxonMobil Chemical Company (2006, 2007) and as a consultant for General Electric Company (2008). His research interests are in the areas of mathematical modeling and optimization of energy systems.

Matthew Rocklin is a Ph.D. computational mathematics student in the Computer Science Department at the University of Chicago. He completed his undergraduate degree in physics and mathematics at the University of California-Berkeley in 2007 and is now focusing on scientific computing.

Sangmin Lee is currently a Ph.D candidate in the Mathematics Department of the Courant Institute at New York University. He earned B.Sc. and M.Sc. degrees in physics from Seoul National University, Korea, in 1999 and 2001, respectively. His research interests lie in the area of numerical methods for stochastic processes.

Mihai Anitescu obtained his Engineer (M.Sc.) Diploma in electrical engineering from the Polytechnic University of Bucharest in 1992 and his Ph.D. degree in applied mathematical and computational sciences from the University of Iowa in 1997. He is currently a computational mathematician in the Mathematics and Computer Science Division at Argonne National Laboratory and a professor in the Department of Statistics at the University of Chicago. He is the author of more than 60 papers in scholarly journals and conference proceedings, on numerical optimization, numerical analysis, computational mathematics and their applications.