Economic Impacts of Advanced Weather Forecasting in Energy System Operations

Victor M. Zavala, Emil M. Constantinescu, and Mihai Anitescu

***Abstract*—We analyze the impacts of adopting advanced weather forecasting systems at different levels of the decision-making hierarchy of the power grid. Using case studies, we show that state-of-the-art numerical weather prediction (NWP) models can provide high-precision forecasts and uncertainty information that can significantly enhance the performance of planning, scheduling, energy management, and feedback control systems. In addition, we assess the forecasting capabilities of the Weather Research and Forecast (WRF) model in several application domains.**

***Index Terms*—weather forecasting, operations, economics, photovoltaics, wind, storage, unit commitment, energy dispatch.**

I. Introduction

Weather conditions strongly influence the energy consumption and performance of industrial and residential facilities. One of the reasons is that energy is used mostly for conditioning large and exposed spaces (e.g., buildings and warehouses) and for generating industrial-scale utilities (e.g., steam and cooling water) that rely on combustion and natural convection phenomena in the presence of atmospheric air. In addition, the next-generation energy portfolio relies on the exploitation of intermittent resources such as solar radiation and wind power.

The operational hierarchy of energy systems involves several layers, including planning and scheduling (or commitment), set-point optimization (or energy management), and feedback control. The weather acts as a dynamic disturbance to all these operational layers. Major questions arise in this context: *What is the economic value of forecasting the weather conditions? How can this information be exploited efficiently?How long and precise should the forecast be?* Intuitively, one would expect that anticipating weather trends allows the operational layers to react more proactively and thus enhances responsiveness and robustness. Forecast trends can be exploited systematically through optimization formulations of the different operational layers. In certain domains where the economic impact of the forecast accuracy is strong, the formulation might also need to incorporate uncertainty information in order to guarantee appropriate performance and robustness. An important industrial example is unit commitment/energy dispatch, where ambient temperature forecasts are used to infer the load. In this operational layer, it is desirable that the spread of the forecast error distribution be as narrow as possible in order to minimize reserves, since these translate into extra costs. In some other domains, however, the economic value of using weather forecasting is entirely system and task dependent and is difficult to assess.

Weather forecasting can be performed by using empirical models (e.g., autoregressive, neural networks, persistence) or advanced numerical weather prediction models (e.g., physics-based). A common problem with empirical models is that their predictive capabilities rely strongly on the presence of persistent trends. In addition, they neglect the presence of coupled, spatiotemporal physical phenomena. Thus, they can lead to inaccurate medium and long-term forecasts and over- or underestimated uncertainty levels [9], [5]. The use of physics-based numerical weather prediction (NWP) models is thus desirable because longer-term forecasts and narrower uncertainty bounds can be obtained. On the other hand, from an operational point of view, 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. This is an important issue because NWP models are expected to be used to make low- and high-frequency operational decisions (time scales from days to minutes).

In this work, we present an integrative study of weather forecasting, uncertainty quantification, and optimization-based operations. We first analyze the impact of increasing the forecast horizon on energy management operations for a multi storage hybrid system and on the temperature control of a building system. With this information, we seek to understand under which conditions weather forecasts are beneficial. We then compare the forecasting capabilities of an empirical modeling approach with those of the Weather Research and Forecast (WRF) model. Here, we seek to demonstrate that WRF can provide accurate and consistent uncertainty bounds for different weather variables in reasonable computational time. However, limitations do exist in some domains. Therefore, we discuss implementation bottlenecks and sources of error excursions, and we propose potential enhancements.

II. Economic Impact of Forecasting

In this section, we analyze two systems to illustrate the economic impact of folding weather forecasts in operations.

A. Multi-Storage Hybrid System

We first consider a photovoltaic system coupled to two storage options described in [14], [16] and sketched in Fig. 3. The first storage option has a large capacity but low round-trip efficiency (hydrogen with 70% efficiency), while the second has a small capacity but high efficiency (battery with 90% efficiency). The operating principle of this system is similar to that of other multi storage systems, such as photovoltaic-compressed air-battery ,or wind-hydrogen-hydrothermal systems [10]. In these systems, it

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

is necessary to decide the best strategy to store the intermittent power input in order to minimize power losses and satisfy a given load. To compute the optimal policies, we formulate an optimal control problem with forecast horizons ranging from 1 hour to 14 days [17]. In Fig. 2, we present the effect of increasing the horizon on the relative operating costs (using a one-year forecast policy as reference). Several conclusions can be drawn from this study: (1) the relative operating costs decay quickly to zero as the horizon is increased; (2) for a purely reactive strategy (1 hr), the relative costs can go as high as 300%; and (3) the close-to-optimal costs can be obtained with a relatively short forecasts (1-14 days). The economic penalty of using a forecast of 1 day is just an increase of 10% in relative costs, whereas the penalty for a forecast of 12 hr goes up to 31%. This implies that a practical forecast horizon should capture the periodicity of the daily radiation. For this system, as the horizon is increased, it is possible to exploit the more efficient battery system to reduce the power losses of the hydrogen storage loop. The magnitude of the cost reductions is problem dependent, but the behavior is expected to be consistent in other storage systems. The major factors affecting the impact of the forecast are the relative efficiency and size of the storage devices.

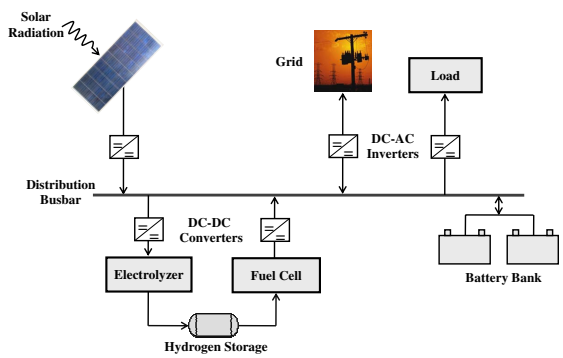

Fig. 1. Schematic representation of multi storage hybrid system.

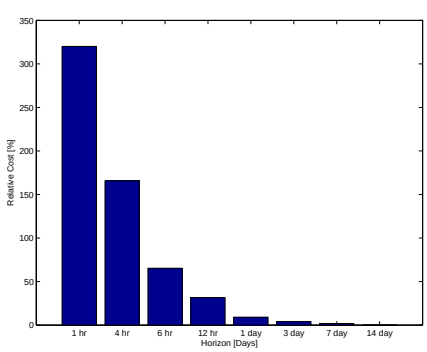

Fig. 2. Effect of forecast horizon on operational costs of multi storage hybrid system.

B. Building System

Commercial buildings are energy-intensive facilities where considerable cost savings can be realized through optimal temperature control strategies. Researchers have found that the

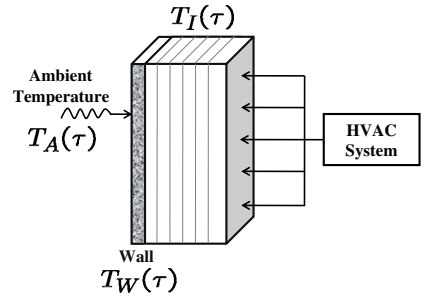

Fig. 3. Schematic representation of building system.

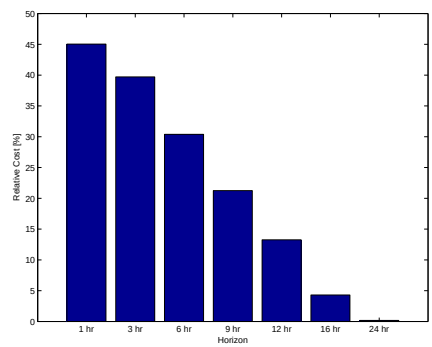

Fig. 4. Effect of forecast horizon on operational costs of building system.

thermal mass of a building can be used for temporal energy storage [2]. With this, it is possible to optimize the temperature set-points trajectories during the day to shift the heating and cooling electricity demands to off-peak hours and thus reduce costs [3]. Since the thermal response of the building can be slow (order of hours), this can be exploited to reduce the on-peak electricity demand the next day. However, the optimal timing at which to start the cooling at night directly depends on the ambient temperature expected the next day. In addition, special care needs to be taken to stay within the thermal comfort zone at all times. To analyze the effect of increasing the forecast horizon of the ambient temperature, we formulate an optimal control problem of the general form presented in Section IV. The problem involves considering the dynamic response of the building internal temperature and of the building wall and trying to find the optimal temperature set-point that minimizes the heating and cooling costs under a given electricity price structure. The ambient temperature enters the model through a boundary condition at the external face of the wall [15].

The relative costs are presented in Fig. 4. As can be seen, for a purely reactive strategy, the relative costs can go as high as 24%. In addition, we observe that a horizon of 1 day is sufficient to achieve the minimum potential costs. The reason is that the thermal mass of the building cannot be used for a long time because there exist losses through the wall. In fact, we found that if the building insulation is enhanced, the costs for the purely reactive strategy increase significantly. On the other hand, when the building is poorly insulated, increasing the forecast horizon does not reduce the costs. In other words, the economic potential of adding forecast information is tightly related to the ability to store energy in the system and to use it during off-peak

times. Another potential benefit of using forecast information is to minimize the number of start ups and shut downs of water chillers and gas furnaces, thereby reducing electricity and fuel costs and enhancing the responsiveness of the HVAC system.

III. Weather Forecasting

In this section, we contrast the forecasting capabilities of empirical and physics-based weather models.

A. Gaussian Process Model

A straightforward forecasting alternative is to use historical measurement data to construct regression models. Several empirical modeling techniques can be used to generate weather forecast trends. An approach that has recently received attention is Gaussian process (GP) modeling [11], [12], [7], [8]. The idea is to construct an autoregressive model by specifying the structure of the covariance matrix rather than the structure of the dynamic model itself, as in traditional system identification techniques such as the Box-Jenkins approach [1]. To illustrate the use of this technique, we construct a forecast model by regressing the future weather variable value (output) χ_{k+1} to the current and previous values (inputs) $\chi_k, ..., \chi_{k-N}$ that can be obtained from weather information data bases. In this case, N is selected long enough to capture the trends of the variable of interest. We define the model inputs as $\mathbf{X}_{[j]} = [\chi_{k-N-j}, ..., \chi_{k-j}]$ and the outputs as $\mathbf{Y}_{[j]} = \chi_{k+1-j}$, and we collect a number of training sets $j = 0, ..., N_{train}$. We assume that the inputs are correlated through an exponential covariance function of the form

$$\mathbf{V}(\mathbf{X}_{[j]}, \mathbf{X}_{[i]}, \eta) := \eta_0 + \eta_1 \cdot \exp\left(-\frac{1}{\eta_2}\|\mathbf{X}_{[j]} - \mathbf{X}_{[i]}\|^2\right), \quad i = 0, ..., N_{train}, j = 0, ..., N_{train},$$

where η_1, η_2, and η_3 are hyperparameters estimated by maximizing the log likelihood function

$$\log p(\mathbf{Y}|\eta) = -\frac{1}{2}\mathbf{Y}^T\mathbf{V}^{-1}(\mathbf{X}, \mathbf{X}, \eta)\mathbf{Y} - \frac{1}{2}\log\det(\mathbf{V}(\mathbf{X}, \mathbf{X}, \eta)).$$

Once the optimal hyperparameters η^* are obtained, we can compute mean predictions $\mathbf{Y}^P$ with associated covariance $\mathbf{V}^P$ at a set of test points $\mathbf{X}^P$. In our context, these are the evolving weather trends. The resulting GP posterior distribution is

$$\begin{aligned}\mathbf{Y}^P &= \mathbf{V}(\mathbf{X}^P, \mathbf{X}, \eta^*)\mathbf{V}^{-1}(\mathbf{X}, \mathbf{X}, \eta^*)\mathbf{Y} \\ \mathbf{V}^P &= \mathbf{V}(\mathbf{X}^P, \mathbf{X}^P, \eta^*) \\ &\quad -\mathbf{V}(\mathbf{X}^P, \mathbf{X}, \eta^*)\mathbf{V}^{-1}(\mathbf{X}, \mathbf{X}, \eta^*)\mathbf{V}(\mathbf{X}, \mathbf{X}^P, \eta^*).\end{aligned}$$

The inverse of the input covariance $\mathbf{V_X} := \mathbf{V}^{-1}(\mathbf{X}, \mathbf{X}, \eta^*)$ (e.g., its factorization) needs to be computed only during the training phase. With this, we can define a conceptual GP model of the form

$$\mathbf{Y}^P = \mathbf{GP}(\mathbf{X}^P, \eta^*, \mathbf{V_X}).$$

Note that at current time t_k, we have measurements to compute only the single-step forecast $\bar{\chi}_{k+1}$. To obtain multi step forecasts, we must propagate the GP predictions recursively. We use the following algorithm,

1) **Forecast mean computation:** For $j = 1, ..., N_F$ do,
   a) Set $\mathbf{X}^P_{[j]} \leftarrow [\chi_{k-N}, \chi_{k-N+1} ..., \chi_k]$
   b) Compute $\mathbf{Y}^P_{[j]} = \mathbf{GP}(\mathbf{X}^P_{[j]}, \eta^*, \mathbf{V_X})$
   c) Drop last measurement, set $\chi_{k+1} \leftarrow \mathbf{Y}^P_{[j]}$, and update $k \leftarrow k+1$
2) **Forecast covariance computation:** Compute self-covariance $\mathbf{V}(\mathbf{X}^P, \mathbf{X}^P, \eta^*)$ and cross-covariance $\mathbf{V}(\mathbf{X}^P, \mathbf{X}, \eta^*)$. Compute forecast covariance $\mathbf{V}^P$ from (1a).

This recursion generates the forecast mean $\mathbf{Y}^P = [\bar{\chi}_{k+1}, ..., \bar{\chi}_{k+N_F}]$ and associated covariance matrix $\mathbf{V}^P$. The resulting distribution $\mathcal{N}(\mathbf{Y}^P, \mathbf{V}^P)$ can be sampled to generate realizations of the future weather trends. Note that the forecasted trends are local (single point in space) and thus do not capture effects at neighboring locations.

B. WRF Model

We now describe the procedures used to compute forecasts using WRF. The WRF model [13] is a state-of-the-art NWP 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. 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 to the 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 similar approach is presented in [15].

1) Ensemble Initialization: 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, that is, simulated atmospheric states reconciled with observations in the data assimilation step. 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 be represented only statistically. We use a normal distribution of the initial conditions to describe the confidence in the knowledge of the initial state of the atmosphere. The distribution is centered on the NARR field at the initial time, the most accurate information available. 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 $\mathcal{N}(x_{\text{NARR}}, \mathbf{V})$:

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

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 is typically localized in space. Therefore, in creating the initial ensemble, one needs to estimate the spatial scales associated with each variable. To obtain these scales, we build correlation matrices of the forecast errors using the WRF model. These forecast errors are estimated by using the NCEP method [6], which is based on starting several simulations staggered in time in such a way that, at any time, two forecasts are available. 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. The covariance and correlation matrices can be approximated by $\mathbf{V} \approx \mathbf{d}\mathbf{d}^{\mathbf{T}}$.

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:

$$x_s^{t_F} = \mathcal{M}_{t_0\to t_F}\left(x_s^{t_0}\right) + \eta_s(t)\,,\ \ s \in \{1\ldots N_S\}\,, \qquad (3)$$

where $x_s^{t_0} \sim \mathcal{N}\left(x_{\text{NARR}}, \mathbf{V}^{t_0}\right)$, $\eta_s \sim \mathcal{N}\left(0, \mathbf{Q}\right)$, 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 [15].

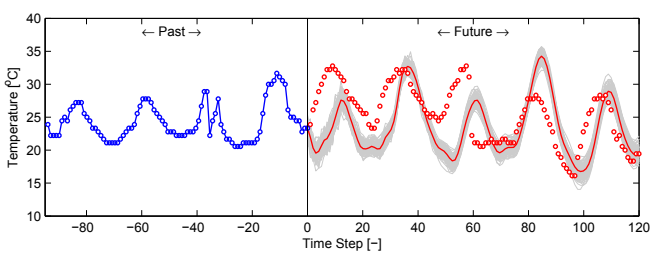

Fig. 5. Ambient temperature observations and GP model realizations for a five-day forecast horizon.

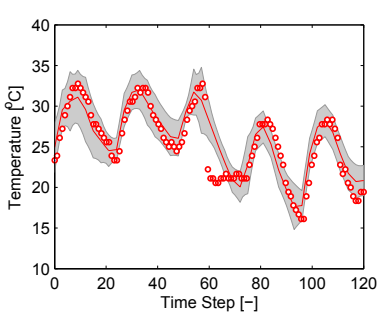

Fig. 6. Ambient temperature observations and WRF model realizations for a five-day forecast horizon.

C. Model Validation

We next validate the forecast information obtained from the WRF and the GP models. An ambient temperature data set at position 40 30'N/80 13'W in the Pittsburgh, PA, area for year 2006 was used. The data were obtained from the National Weather Service Office. To assess the forecasting capabilities of the GP model, we used a total of 120 training sets and we set N to 24. We consider a forecast horizon of 5 days. In Fig. 5, we present the forecast mean and 100 realizations drawn from the predictive distribution. We can see that the GP model is able to capture the periodicity of the trends. However, the mean drifts away from the true realizations and the uncertainty bounds are not able to encapsulate the actual realizations. Since the distribution is inconsistent, it is not useful from an operational point of view. Note that the temperature follows variations as a result of spatial interactions and long-term metereological phenomena. This situation might explain the lack of stronger periodic trends. This might also explain the poor quality of the forecasts. Nevertheless, we emphasize that empirical models are valuable for high-frequency operations (on the order of seconds or minutes). We now validate the forecast and uncertainty

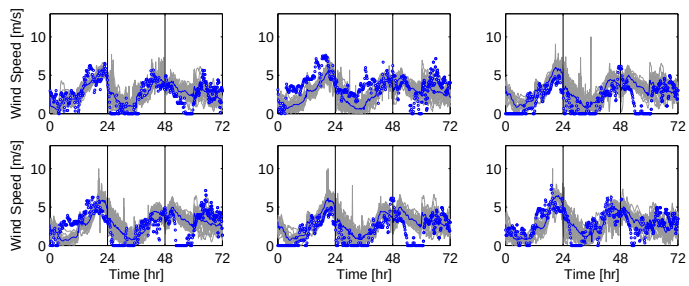

Fig. 7. Wind speed realizations for 6 wind-farm locations in Illinois at 10 m and observations (dots) at nearest meteorological stations. Vertical lines represent beginning of day.

information of the WRF model. In Fig. 6, we show the predictive distribution for the temperature constructed using 30 ensemble members and a forecast horizon of 5 days. Note that the forecast errors are small ($\pm 5^oC$) and the uncertainty bounds enclose the true observations tightly. We also present validation results of wind speed forecasts generated by WRF at 6 active wind-farms in the state of Illinois to analyze their accuracy and correlation structure. The wind speed fields at 10 meters above the ground for three consecutive days are presented in Fig. 7. We note that the WRF realizations are able to capture the general trends of the actual observations and to encapsulate the observations with few

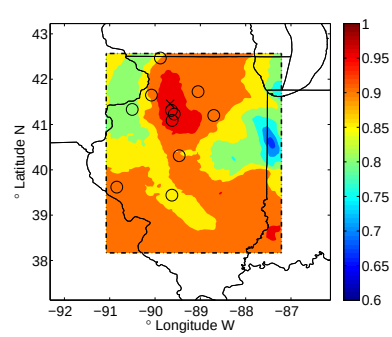

Fig. 8. The spatial correlation for the wind field for wind farm #8 on June 5, 1:50 AM, denoted by "X." The circle markers denote the other wind farms in Illinois.

exceptions. The largest differences are observed at the beginning of the third day. The wind speed trends are clearly more difficult to predict than temperature trends. In Fig. 8 we show the spatial correlations of the wind speed for a particular wind farm 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. This observation is confirmed by comparing Figs. 7 and 8. Here, we can see that the wind speed realizations for wind farms #2, 3, and 4 are strongly correlated, as predicted by the correlation mapping.

In this study we used version 3.1 of WRF [13]. The ensemble approach taken for estimating the uncertainty in the weather system is highly parallelizable because each realization 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-grain level consists in the parallelization of each sample. Our running times indicate that around 32 CPUs are sufficient to generate updated forecasts with WRF for intra day operations (e.g., unit commitment). The times also suggest that, in order to generate updated forecasts for hourly operations (e.g., energy management), one would need about 500 CPUs.

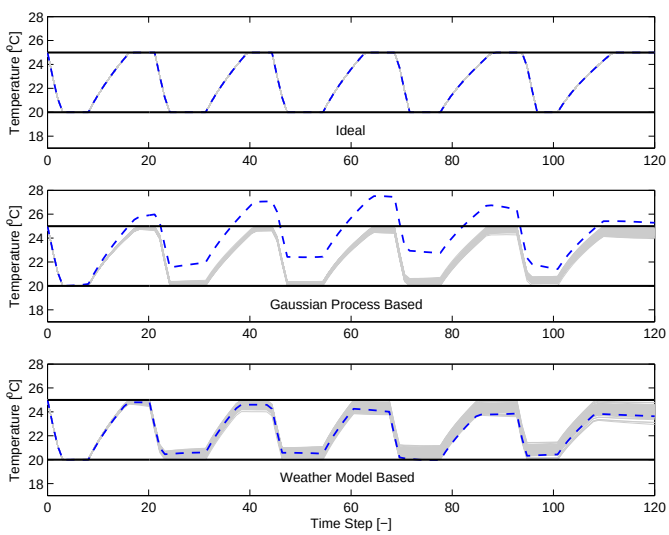

Fig. 9. Operating policies with different forecasts. Thermal comfort zone is highlighted by thick solid lines, predicted temperatures are gray lines, and actual realizations are dashed lines.

IV. Integration of Forecasting and Operations

In this section we discuss how forecast accuracy and uncertainty information affect the performance of different operational layers.

A. Building Temperature Control

As discussed in Section II, optimal set-points for internal building temperatures can be obtained by solving an optimal control problem. We can extend this formulation to consider multiple future realizations of the ambient temperature. The problem can be cast as a sample-average stochastic optimal control problem:

$$\min_{u(\tau)} \frac{1}{N_s}\sum_{s=1}^{N_s}\left[\int_{t_k}^{t_k+T}\varphi(z_s(\tau),y_s(\tau),u(\tau),\chi_s(\tau))d\tau\right] \tag{4a}$$

$$\text{s.t.}\ \left.\begin{aligned}\frac{dz_s}{d\tau} &= \mathbf{f}\left(z_s(\tau),y_s(\tau),u(\tau),\chi_s(\tau)\right)\\ 0 &= \mathbf{g}(z_s(\tau),y_s(\tau),u(\tau),\chi_s(\tau))\\ 0 &\geq \mathbf{h}(z_s(\tau),y_s(\tau),u(\tau),\chi_s(\tau))\end{aligned}\right\}\ \tau\in[t_k,t_k+T] \tag{4b}$$

$$z_s(t_k) = x_k, \quad s\in\mathcal{S}. \tag{4c}$$

The objective function is the future average cost; the constraints include a differential-algebraic equation system describing the building dynamics and a set of operational constraints. The dynamic model is defined over each possible realization of the weather conditions, which are denoted by $\chi_s(\tau),\ s\in\mathcal{S}$, and presented in Figs. 5 and 6. We solve this problem using realizations obtained from the GP and WRF model distributions. The optimal temperature set-point profiles are presented in Fig. 9. In the top graph, we present the profile for the ideal strategy where perfect forecast information is assumed. Since there is no uncertainty, the predicted profile matches the actual realization. Note that the optimal profile hits continuously the bounds of the comfort zone, as it tries to take advantage of the on-peak and off-peak electricity rates to minimize costs. In the middle graph, we present the optimal profiles obtained using forecast information from the GP model. The gray lines are the predicted realizations of the dynamic model in the stochastic optimal control formulation. Since the uncertainty structure provided by GP is not able to capture the ambient temperature observations, the actual realization of the internal building temperature goes outside the comfort zone. In the bottom graph, we see that the use of the WRF model realizations results in an temperature trajectory that stays within the comfort zone at all times. Note that, since the comfort zone is very narrow ($\approx 5^o$C), high-precision forecast information is needed to realize the economic benefits.

B. Unit Commitment and Energy Dispatch

We now analyze the accuracy of wind speed forecasts obtained with WRF in unit commitment/energy dispatch operations with large adoption levels of wind power. The system specifications are provided in [4]. The sample-average stochastic formulation is

$$\min_{p_{s,j,k},\overline{p}_{s,j,k},\nu_{j,k}} \frac{1}{N_S}\sum_{s\in\mathcal{S}}\left(\sum_{j\in\mathcal{N}}\sum_{k\in\mathcal{T}} c^p_{s,j,k} + c^u_{j,k} + c^d_{j,k}\right) \tag{5a}$$

$$\text{s.t.}\ \sum_{j\in\mathcal{N}} p_{s,j,k} + \sum_{j\in\mathcal{N}_{wind}} p^{wind}_{s,j,k} = D_k,\ s\in\mathcal{S}, k\in\mathcal{T} \tag{5b}$$

$$\sum_{j\in\mathcal{N}} \overline{p}_{s,j,k} + \sum_{j\in\mathcal{N}_{wind}} p^{wind}_{s,j,k} \geq D_k + R_k,\ s\in\mathcal{S}, k\in\mathcal{T} \tag{5c}$$

$$p_{s,j,k}\in\Pi_{s,j,k},\ s\in\mathcal{S}, j\in\mathcal{N}, k\in\mathcal{T}. \tag{5d}$$

Here, $p_{s,j,k}$ are the power outputs of the thermal units, $p^{wind}_{s,j,k}$ are the wind power outputs of a set of turbines, D_k is the demand, R_k are the reserves, and $\Pi_{s,j,k}$ is a set of feasible

outputs that implicitly takes into account ramp, start-up, and shutdown constraints. In Fig. 10, we present the profiles of total aggregated demand, thermal power, and wind power for three days of operation. The wind power profiles were generated by using the wind speed realizations presented in Fig. 7. Note that the aggregated wind power profile 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%. In 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. However, this result suggests that a high frequency and adaptive inflation/resampling procedure is needed in WRF. This clearly points out the need for more application-oriented forecasting capabilities. Moreover, we 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.

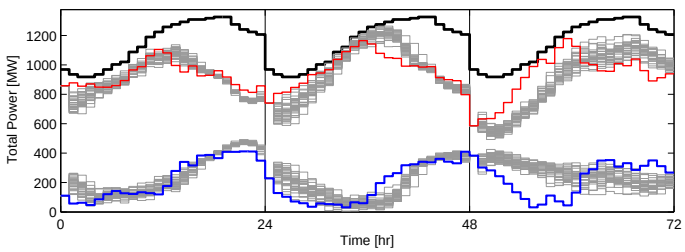

Fig. 10. Closed-loop total power profiles obtained with stochastic UC formulation. Top thick line is demand profile, medium thick line is the implemented thermal profile, gray lines are planned realizations at beginning of each day, bottom thick line is actual total wind power, and the adjacent gray lines are forecast profiles.

V. Conclusions

We have presented an integrative study of weather forecasting, uncertainty quantification, and optimization-based operations to analyze the economic impacts of adopting advanced weather forecasting systems at different operational levels of the power grid. Our studies suggest that costs and power losses can be reduced by incorporating accurate forecasts. In addition, our model validation studies indicate that WRF forecasts significantly outperform those obtained with empirical models, specially for medium- and long-term horizons. As future work, we are interested in developing application-oriented forecasts to minimize computational limitations. In addition, targeted data assimilation and adaptive inflation resampling procedures are necessary to correct forecasts at higher frequencies.

Acknowledgments

This work was supported by the Department of Energy, under Contract No. DE-AC02-06CH11357.

References

[1] G. E. P. Box, G. Jenkins, and G. Reinsel, *Time Series Analysis: Forecasting and Control*, Prentice-Hall, New Jersey, 1994.

[2] J. E. Braun, *Reducing energy costs and peak electricty demand through optimal control of building thermal storage*, ASHRAE Transactions, 96 (1990), pp. 876–888.

[3] J. E. Braun, K. W. Montgomery, and N. Chaturvedi, *Evaluating the performance of building thermal mass control strategies*, HVAC&Research, 7 (2001), pp. 403–428.

[4] E. M. Constantinescu, V. M. 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, 2009.

[5] 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.

[6] 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.

[7] H. K. Lee, D. M. Higdon, C. A. Calder, and C. H. Holloman, *Efficient models for correlated data via convolutions of intrinsic processes*, Statistical Modeling, 5 (2005), pp. 53–74.

[8] C. Paciorek and M. Schervish, *Nonstationary covariance functions for Gaussian process regression*, in Advances in Neural Information Processing Systems 16: Proceedings of the 2003 Conference, Bradford Book, 2004, p. 273.

[9] 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.

[10] 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.

[11] C. E. Rasmussen and C. K. Williams, *Gaussian Processes for Machine Learning*, MIT Press, Cambridge, 2006.

[12] S. K. Sahu, S. Yip, and D. M. Holland, *Improved space-time forecasting of next day ozone concentrations in the eastern US*, Atmospheric Environment, 43 (2009), pp. 494 – 501.

[13] 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 475+ STR, NCAR, 2008.

[14] O. Ulleberg, *The importance of control strategies in PV-hydrogen systems*, Solar Energy, 76 (2004), pp. 323–329.

[15] V. Zavala, E. Constantinescu, T. Krause, and M. Anitescu, *On-line economic optimization of energy systems using weather forecast information*, Journal of Process Control, (2009).

[16] V. M. Zavala, M. Anitescu, and T. Krause, *On the Optimal On-Line Management of Photovoltaic-Hydrogen Hybrid Energy Systems*, in Proceedings of 10th International Symposium on Process Systems Engineering, R. M. de Brito Alves, C. A. O. do Nascimento, and E. C. B. Jr., eds., 2009.

[17] V. M. Zavala, E. M. Constantinescu, and M. Anitescu, *Weather forecast-based optimization of integrated energy systems*, Tech. Report ANL/MCS-TM-305, Argonne National Laboratory, 2009.

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.

Mihai Anitescu has obtained his Engineer (M.Sc.) Diploma in electrical engineering from the Polytechnic University of Bucharest (1992) and his Ph.D. degree in applied mathematical and computational sciences from the University of Iowa (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.

The submitted manuscript has been created by the University of Chicago as Operator of Argonne National Laboratory ("Argonne") under Contract No. DE-AC02-06CH11357 with the U.S. Department of Energy. The U.S. Government retains for itself, and others acting on its behalf, a paid-up, nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.