

On-Line Economic Optimization of Energy Systems Using Weather Forecast Information

Victor M. Zavala,^{†§*} Emil M. Constantinescu,[†]
Theodore Krause,[§] and Mihai Anitescu[†]

[†]Mathematics and Computer Science Division

[§]Chemical Sciences and Engineering Division

Argonne National Laboratory, Argonne, IL, 60439 United States

Abstract

We establish an on-line optimization framework to exploit weather forecast information in the operation of energy systems. We argue that anticipating the weather conditions can lead to more *proactive* and cost-effective operations. The framework is based on the solution of a stochastic dynamic real-time optimization (D-RTO) problem incorporating forecasts generated from a state-of-the-art weather prediction model. The necessary uncertainty information is extracted from the weather model using an ensemble approach. The accuracy of the forecast trends and uncertainty bounds are validated using real meteorological data. We present a numerical simulation study in a building system to demonstrate the developments.

1 Introduction

The economic performance of industrial systems is strongly affected by dynamic disturbances evolving at different time scales. These include input flows, product demands, energy prices, weather conditions, among others. In order to manage these disturbances, the operational decisions are decomposed in a hierarchical manner. The top decision-making level is the supervisory or economic optimization layer which adjusts the set-points as low-frequency disturbances evolve in time. The lower decision-making level is the control level that rejects high-frequency disturbances in order to keep the process close to economic optimal set-points. A widely-used supervisory operation strategy is real-time optimization (RTO) [35]. RTO makes use of a steady-state rigorous process model to determine the optimal set-points that maximize the system profit under the *current* disturbance information. An advantage of this strategy is that it can be seen as a closed-loop optimizer

*Address all correspondence to this author. Argonne National Laboratory, Mathematics and Computer Science Division, 9700 S. Cass Avenue, Argonne, IL 60439, USA. email:vzavala@mcs.anl.gov

that *rejects* disturbances affecting profit [45]. Nevertheless, an important limitation of RTO is that it is entirely *reactive*, in the sense that only current disturbances are taken into account. In other words, RTO neglects the fact that disturbances follow trends that can be *exploited* to obtain more efficient operating policies. An alternative to overcome this limitation is to incorporate disturbance *forecast* information in a dynamic real-time optimization (D-RTO) formulation [29]. The exploitation of disturbance trends can add *proactiveness* to the operational decisions which can be beneficial both from an economic and a control perspective. In particular, we claim that the capabilities of D-RTO can be greatly expanded through the incorporation of *weather* forecast information. This can be particularly critical in the operation of energy production systems since their performance is strongly affected by the evolution of the weather conditions. Polygeneration energy systems [11], power plants [25], wind farms [24], photovoltaic systems [50; 54], and building climate control [7; 40] represent some important applications in the energy sector. Some other potential industrial applications are the optimization of utility systems in chemical complexes and the optimization of wastewater treatment plants [9].

Weather forecast information has been used in diverse operational studies. In particular, it has been widely used for planning, scheduling, and unit commitment tasks in power grid operations [53; 19]. At these higher decision-making levels, weather information is exploited indirectly by mapping it to economic variables such as user power demands (e.g., electricity and heating/cooling needs). However, at lower economic optimization and control levels, trends of ambient temperature, wind speed, solar radiation, and humidity can be exploited directly through the rigorous process model.

The evolution of the weather conditions is affected by complex physico-chemical phenomena that are difficult to predict. This becomes relevant if the economic performance is dictated by the weather conditions (e.g., wind farm) or if the system is subject to tight operational constraints (e.g., comfort zone in building systems). In such cases, incorporating forecast uncertainty in the D-RTO formulation is important. This can be done through suitable stochastic or robust optimization formulations, however, some important obstacles arise. First, obtaining detailed uncertainty information validated with real data (i.e., probability distributions) is often difficult. For instance, covariance information and high-resolution (i.e., minutes, hours) weather trends might be required. This detailed information can be expensive or impractical to obtain from commercial weather prediction companies. Second, stochastic optimization problems are infinite-dimensional and special techniques are needed for their solution. Alternatives to solve stochastic dynamic optimization problems include, among others, dynamic programming [4], Taylor series expansions [3; 37], and polynomial chaos expansions [1]. Certain restricted classes of chance-constrained [39; 40] and minmax [34; 17] formulations can also be reformulated and solved using standard optimization techniques. A detailed review of available formulations and solution methods is beyond the scope of this paper. However, we emphasize that a practical obstacle commonly encountered with the aforementioned approaches is that they can only handle restricted problem classes.

In this work, we establish an on-line optimization framework able to exploit detailed weather forecast information. In particular, we extend the capabilities of a numerical

weather prediction model to provide detailed uncertainty information and derive a stochastic D-RTO formulation able to exploit this information. The uncertainty of the weather prediction model is quantified using an ensemble approach. We propose to use a sample-average approximation (SAA) to solve the associated problems. We claim that the SAA approach is attractive from an industrial perspective because it can be implemented quite easily, it can accommodate large-scale models, and can handle general stochastic formulations with restricted uncertainty information. We contrast the forecast capabilities of the detailed weather prediction model against those of empirical autoregressive models. Both forecast models are validated using real meteorological data. We demonstrate that the use of efficient uncertainty quantification techniques is critical to obtain consistent uncertainty bounds and appropriate performance of the stochastic optimizer. We present simulation studies on a building system to illustrate the developments. To the best of our knowledge, this is the first report proposing to integrate stochastic optimization formulations and uncertainty quantification techniques for weather forecasting in order to optimize the economic performance of energy systems.

The paper is organized as follows. In the next section, we establish the stochastic D-RTO formulation. In Section 3, we present techniques to quantify the uncertainty of weather forecasts. A simulation case study is presented in Section 4 and the paper closes with general conclusions and directions for future work.

2 Stochastic Optimization

In this section, we derive a general D-RTO formulation and discuss extensions to consider stochastic disturbance information. We begin by considering a differential-algebraic equation (DAE) model of the form

$$\frac{dz}{d\tau} = \mathbf{f}(z(\tau), y(\tau), u(\tau), \chi(\tau)) \quad (1a)$$

$$0 = \mathbf{g}(z(\tau), y(\tau), u(\tau), \chi(\tau)) \quad (1b)$$

$$z(0) = x_k, \quad (1c)$$

where τ is the *model* time dimension and t_k is the current time in the *real* system. Variables $z(\tau)$ are differential states, $y(\tau)$ are algebraic states, $u(\tau)$ are the controls or manipulated variables, and $\chi(\tau)$ are the exogenous disturbances. In this context, the term *exogenous* refers to the fact that the disturbances are not affected by the system variables (e.g., energy prices). The differential equations (1a) represent conservation equations (energy, mass, and momentum), while the algebraic equations (1b) represent consistency conditions and expressions to calculate physicochemical properties. The initial conditions at time t_k are given by the current state of the system x_k .

The disturbance trajectory $\chi(t)$, $t \in [t_k, t_k + T]$ is random with unknown probability distribution \mathcal{P}_k . However, we assume that this distribution can be *approximated* using a stochastic forecast model (see Section 3). For instance, we can assume that the forecast model provides a predictive mean $\bar{\chi}(\tau)$ and that the associated forecast errors follow a Gaussian distribution. With this, we can approximate \mathcal{P}_k using $\mathcal{N}(\bar{\chi}(\tau), \mathbf{V}(\tau))$, where

$\mathbf{V}(\tau)$ is the covariance matrix. In this case, a fixed probability level of \mathcal{P}_k defines an ellipsoidal region Ω_k of the form,

$$\Omega_k := \{z \mid (z - \bar{\chi}(\tau))^T \mathbf{V}^{-1}(\tau)(z - \bar{\chi}(\tau)) \leq \alpha\}, \quad (2)$$

This region is sketched in Figure 1. Under the Gaussian assumption for the forecast errors, all that is needed to represent the \mathcal{P}_k is the predictive mean and the covariance matrix. However, we emphasize that the proposed structure of the probability distribution is a modeling assumption and hence might not be accurate. Nevertheless, from a practical point of view, *what we seek is that the approximate probability distribution is able to encapsulate the true disturbance realizations and that it has a physically meaningful structure.* To exploit the statistical information at hand, we formulate a stochastic




Figure 1: Schematic representation of ellipsoidal uncertainty region.

dynamic optimization problem of the form

$$\min_{u(\tau)} \mathbf{E}_{\chi(\tau)} \left[\int_{t_k}^{t_k+T} \varphi(z(\tau), y(\tau), u(\tau), \chi(\tau)) d\tau \right] \quad (3a)$$

$$\text{s.t. } \left. \begin{aligned} \frac{dz}{d\tau} &= \mathbf{f}(z(\tau), y(\tau), u(\tau), \chi(\tau)) \\ 0 &= \mathbf{g}(z(\tau), y(\tau), u(\tau), \chi(\tau)) \\ 0 &\geq \mathbf{h}(z(\tau), y(\tau), u(\tau), \chi(\tau)) \end{aligned} \right\} \tau \in [t_k, t_k + T], \quad \forall \chi(\tau) \quad (3b)$$

$$z(t_k) = x_k, \quad (3c)$$

where symbol $\mathbf{E}[\cdot]$ denotes the expectation operator with associated distribution \mathcal{P}_k . From the solution of this problem, we obtain the state and control trajectories $z^*(\tau), y^*(\tau), u^*(\tau), \tau \in [t_k, t_k + T]$ that can be sent to a lower-level controller as set-points. At the next time step t_{k+1} , we obtain the updated state of the system x_{k+1} and the updated forecast distribution \mathcal{P}_{k+1} that we use to solve the next stochastic problem (3).

Note that the weather trends are treated as parameters in the stochastic optimization formulation. In other words, the stochastic differential equations representing the weather dynamics are not introduced considered here. This is justified by the fact that the weather trends act as exogenous disturbances (i.e., the weather is not affected by the system). The weather dynamics will be presented in Section 3. In problem (3) we have assumed that

the expectation of the objective function distribution is an adequate measure of the performance of the system. However, this need not be the case. For instance, we could also choose the mean-risk approach of Markowitz where we seek to minimize simultaneously the mean and the variance of the cost distribution. In a stochastic optimization framework, the structure of the cost function becomes a design task and it is entirely problem dependent. Note also that the proposed stochastic formulation assumes that no recourse exists in the future, as in a dynamic programming approach [4] or in closed-loop MPC approaches [40]. It is well-known that introducing future recourse can reduce the conservativeness of the solution. Finally, we do not consider uncertainty on the system state x_k and model errors. Here, we use a rather simple stochastic optimization formulation in order to illustrate the main points of the proposed strategy.

The stochastic formulation presented in (3) is computationally demanding since it is infinite-dimensional in $\chi(\tau)$. In addition, the number of disturbances is expected to be large as well. To solve the stochastic optimization problem, we use a sample-average approximation (SAA) approach. The idea is to use Monte Carlo sampling to obtain N_s independent samples from the distribution \mathcal{P}_k . This gives the set of disturbance realizations $\{\chi_1(\tau), \chi_2(\tau), \dots, \chi_{N_s}(\tau)\}$. With this, the approximate stochastic problem becomes,

$$\min_{u(\tau)} \frac{1}{N_s} \sum_{j=1}^{N_s} \left[\int_{t_k}^{t_k+T} \varphi(z_j(\tau), y_j(\tau), u(\tau), \chi_j(\tau)) d\tau \right] \quad (4a)$$

$$\text{s.t.} \left. \begin{aligned} \frac{dz_j}{d\tau} &= \mathbf{f}(z_j(\tau), y_j(\tau), u(\tau), \chi_j(\tau)) \\ 0 &= \mathbf{g}(z_j(\tau), y_j(\tau), u(\tau), \chi_j(\tau)) \\ 0 &\geq \mathbf{h}(z_j(\tau), y_j(\tau), u(\tau), \chi_j(\tau)) \end{aligned} \right\} \tau \in [t_k, t_k + T], \quad j = 1, \dots, N_s \quad (4b)$$

$$z_j(t_k) = x_k. \quad (4c)$$

In this formulation, all the variables become a function of the particular disturbance realization except the controls, which are decision variables. One of the key advantages of the SAA approach is that it is straightforward to implement. Moreover, it is particularly suitable for large-scale systems because it gives rise to highly structured problems [33; 28]. In addition, it allows the solution of general stochastic formulations in a unified manner. The theoretical properties of the SAA approach have been studied in the context of nonlinear programming. For instance, it has been shown that solutions of the SAA problem converge at an exponential rate with the number of samples to the solution of the stochastic counterpart [47; 2]. In addition, the number of scenarios N_s used in the SAA problem can be related to the probability of satisfying the constraints. With this, equivalence with chance constrained formulations can be established [39]. Although no formal convergence results exist in the context of dynamic optimization problems, we can expect that the available convergence guarantees can be used under mild assumptions. For instance, if we discretize problem (3) in time, we can argue that as long as we have a convergent and well-posed time-discretization scheme, the SAA approach will converge to the stochastic counterpart under standard regularity assumptions [47; 2]. This is an important theoretical question and will not be pursued here. Another crucial advantage of

the SAA approach is that disturbance realizations can be incorporated in the optimization formulation *without even knowing their distribution*. This permits to handle random variables with non-Gaussian distributions and to accommodate disturbance realizations directly in the formulations (i.e., the covariance matrix is not needed). As we will see in the next section, this is important if a detailed weather prediction model is used since the amounts of data to be handled in the optimization framework can become extremely large.

The prototypical closed-loop D-RTO strategy based on SAA is as follows:

1. **Obtain current state and forecast:** At time t_k , obtain current state x_k , and disturbance realizations $\{\chi_1(\tau), \chi_2(\tau), \dots, \chi_{N_s}(\tau)\}$, $\tau \in [t_k, t_k + T]$ from forecast capability.
2. **Compute set-points:** Solve stochastic optimization problem (4). Send optimal set-points $z^*(\tau), y^*(\tau), u^*(\tau), \tau \in [t_k, t_k + T]$ to low-level control layer.
3. **Update:** At $t_k + \Delta$, set $k \leftarrow k + 1$, and repeat process.

Here, Δ is the set-point update period. If \mathcal{P}_k is approximated using a Gaussian distribution, the forecasting capability can also communicate the mean $\bar{\chi}(\tau)$ and covariance matrix $\mathbf{V}(\tau), \tau \in [t_k, t_k + T]$ instead of the disturbance realizations.

3 Uncertainty Quantification

From the previous section, it is clear that different techniques can be used to solve stochastic dynamic optimization problems. However, we emphasize that regardless of the solution approach used, a perhaps more important (and often overlooked) issue in stochastic optimization studies is the need of *consistent* uncertainty information. For instance, constraint satisfaction cannot be enforced appropriately with any solution approach if the bounds of the uncertainty region do not encapsulate the actual realizations. In addition, the conservativeness of the solution is directly related to the uncertainty bounds. Consequently, several questions arise: Can we get accurate and consistent forecast information? What techniques can be used to quantify uncertainty? How does the accuracy of the uncertainty information manifest in the solution of the stochastic optimization problem? In this section, we present two techniques to generate forecasts and uncertainty information that can be exploited by stochastic optimization formulations.

3.1 Numerical Weather Prediction Model

We first derive an ensemble data assimilation approach based on a detailed numerical weather prediction (NWP) model. In particular, we discuss how to capture the uncertainty of the temperature field.

Major weather prediction centers, such as the European Centre for Medium- Range Weather Forecasts (ECMWF) and the U.S. National Centers for Environmental Prediction

(NCEP), are capable of producing high-precision weather forecasts several times a day. Advances in this area are due to improved models of the atmosphere, greater availability of atmospheric data, increased computational power, and the continued improvement of state estimation algorithms. If the state of the atmosphere (temperatures, pressures, humidities, wind speeds and directions) were known exactly at a given time, a forecast could be obtained by integrating the atmospheric model equations forward in time. In practice, however, the state of the atmosphere is known only indirectly through observations that are distributed non-uniformly in space and time and subject to error [22]. Therefore, the model must be first reconciled to the most recent observations. This state estimation problem is called in the weather forecast literature as the *data assimilation* problem. Traditional assimilation techniques such as Kalman filtering [21; 8] and 4-D Var (moving horizon estimation) [16] are used internally in the models for this [13; 14; 15; 41; 26; 27]. The objective function is derived from Bayesian or maximum likelihood principles (e.g., weighted least-squares) and includes a term that penalizes deviation from a prior state (e.g., initial conditions). Once the current state is obtained from the solution of the assimilation problem, it is possible to integrate the atmospheric model equations forward in time to produce a forecast.

The forecast uncertainty can be constructed from the *posterior* distribution of the current state obtained from the data assimilation step. This distribution can then be sampled and evolved through the NWP model dynamics. The resulting trajectories can then be assembled to obtain an approximation of the forecast covariance matrix. This procedure is explained in Section 3.1. In Section 3.1.2, we present a procedure to approximate the *posterior* covariance of the spatial temperature field using the assimilated state of the NWP model.

The Weather Research and Forecasting (WRF) model is a state-of-the-art mesoscale numerical weather prediction system designed to serve both operational forecasting and atmospheric research needs [49]. We use the current version of this model, WRF 3.1, with the default settings for the forecast and uncertainty estimation on temperature fields. The data used in the WRF model in this work corresponds to North American Regional Reanalysis data set that covers 160W-20W, 10N-80N, with a resolution of 10 minutes of a degree. There are 29 pressure levels (1000-100 hPa, excluding the surface) and a three-hour output frequency. The time period under consideration ranges from August 1 to August 30, 2006, [<http://dss.ucar.edu/pub/narr>]. This data set includes meteorological fields such as temperature, wind, and humidity.

3.1.1 Ensemble Approach to Uncertainty Quantification

The dimensions of the state vector in the weather model with a coarse spatial discretization are $\mathcal{O}(10^6) - \mathcal{O}(10^8)$. Therefore, the state covariance matrices are extremely large (grows with the square of the number of states). Hence, in practice, these matrices need to be approximated with a reduced model [12; 13; 14] or with an ensemble of realizations [43; 31; 46]. In this work, we prefer to use the ensemble approach because it can be implemented by using the WRF model as a black box. If the dimension of a random variable x is defined as n , a given covariance matrix $\mathbf{V} \in \mathbb{R}^{n \times n}$ can be approximated by

an ensemble of m realizations x_i , $1 \leq i \leq m$ as,

$$\begin{aligned}\mathbf{V} &:= \frac{1}{m-1} \sum_{i=1}^m (x_i - \bar{x})(x_i - \bar{x})^T \approx \mathbf{E} \left[(x - \bar{x})(x - \bar{x})^T \right], \\ \bar{x} &:= \frac{1}{m} \sum_{i=1}^m x_i \approx \mathbf{E}[x].\end{aligned}\tag{5a}$$

One also has that,

$$\mathbf{V} = D^{\frac{1}{2}} C D^{\frac{1}{2}}, \quad C_{i,j} = \frac{\mathbf{V}_{i,j}}{\sqrt{D_{i,i}} \sqrt{D_{j,j}}} = \frac{\mathbf{V}_{i,j}}{\sigma_i \sigma_j} = \frac{\sigma_{i,j}^2}{\sigma_i \sigma_j}, \quad 1 \leq i, j \leq n,$$

where C is the correlation matrix and D is a diagonal matrix holding the local variances ($D_{i,i} = \sigma_i^2$). In the context of the stochastic optimization framework of Section 2, x represents the future trajectory exogenous states or disturbances $\chi(\tau)$ with mean $\bar{x} := \bar{\chi}(\tau)$ and covariance $\mathbf{V} := \mathbf{V}(\tau)$. Note that the disturbance at a particular point in time t_k represents a three-dimensional spatial field $\chi(t_k, x, y, z)$. In the following, we simplify the notation by eliminating the explicit dependence on the space dimensions. Using this representation, we now discuss how to approximate the forecast covariance.

At current time t_k , the exact state of the atmosphere is not known exactly and therefore the state of the weather model is described through random variables. The errors in atmospheric models are assumed to stem from many independent sources, and thus it is common, realistic, due to the central limit theorem [18], and convenient to consider them as having a normal distribution [30; 26; 36; 20]. The mean of this *posterior* distribution is the assimilated state of the WRF model, which has been reconciled to past measurement data. We denote the *true* state of the weather model at time t_k as $\chi_{true}(t_k)$. Since the weather model is not perfect, the true state at t_{k+1} is given by the stochastic discrete-time model of the form,

$$\chi_{true}(t_{k+1}) = \mathcal{M}(\chi_{true}(t_k)) + \eta(t_k),\tag{6}$$

where \mathcal{M} is the WRF model and $\eta(\cdot)$ represents the model errors that are assumed to be unbiased with covariance Q , $\eta \in \mathcal{N}(0, Q)$. Since the current state is not known exactly, the numerical prediction at time t_{k+1} , $\chi(t_{k+1})$, is obtained from the model evolution of the *believed* state (true solution perturbed with errors). This is represented by a set of unbiased random variables $\varepsilon(t_k)$ with Gaussian distribution $\mathcal{N}(0, \mathbf{V}(t_k))$. With this, we can express the future believed state as

$$\chi(t_{k+1}) = \mathcal{M}(\chi_{true}(t_k) + \varepsilon(t_k)).\tag{7}$$

Since $\chi(t_{k+1})$ becomes a random variable, we define its covariance matrix as $\mathbf{V}(t_{k+1})$. The matrix is given by,

$$\mathbf{V}(t_{k+1}) = \mathbf{E} \left[(\chi(t_{k+1}) - \chi_{true}(t_{k+1})) (\chi(t_{k+1}) - \chi_{true}(t_{k+1}))^T \right]\tag{8a}$$

$$\begin{aligned}&= \mathbf{E} \left[(\mathcal{M}(\chi_{true}(t_k) + \varepsilon(t_k)) - (\mathcal{M}(\chi_{true}(t_k)) + \eta(t_k))) \right. \\ &\quad \left. \cdot (\mathcal{M}(\chi_{true}(t_k) + \varepsilon(t_k)) - (\mathcal{M}(\chi_{true}(t_k)) + \eta(t_k)))^T \right].\end{aligned}\tag{8b}$$

Note that this formula reduces to the Kalman filter covariance update under certain special conditions. To illustrate this, we first assume that the initial condition errors $\varepsilon(t_k)$ and model errors $\eta(t_k)$ are uncorrelated. Consequently,

$$\mathbf{E} [(\mathcal{M}(\chi_{true}(t_k) + \varepsilon(t_k))) \eta(t_k)^T] = 0.$$

Moreover, if we assume that the error growth is well approximated by a linearized model, then

$$\mathcal{M}(\chi_{true}(t_k) + \varepsilon(t_k)) - \mathcal{M}(\chi_{true}(t_k)) = M \cdot \varepsilon(t_k).$$

where $M = \frac{d\mathcal{M}}{dy}$. It follows that (8a) becomes

$$\begin{aligned} \mathbf{V}(t_{k+1}) &\approx \mathbf{E} \left[(M \cdot \varepsilon(t_k) + \eta(t_k)) (M \cdot \varepsilon(t_k) + \eta(t_k))^T \right], \\ &= \mathbf{E} [M \varepsilon(t_k) \varepsilon(t_k)^T M^T] + \mathbf{E} [M \cdot \varepsilon(t_k) \eta(t_k)^T] + \mathbf{E} [\eta(t_k) \varepsilon(t_k)^T M^T] + \mathbf{E} [\eta(t_k) \eta(t_k)^T], \\ &= M \mathbf{V}(t_k) M^T + Q. \end{aligned} \quad (9)$$

Equation (9) represents a linear approximation of the exact error covariance. It is well known that such an approximation can deviate significantly in highly nonlinear models. The ensemble approach proposed in this work propagates the uncertainties in the *current* state field through the nonlinear WRF model according to (7). The covariance matrix is approximated by using an ensemble of realizations generated by sampling the posterior distribution $\mathcal{N}(0, \mathbf{V}(t_k))$. We obtain multi-step trajectories by recursive model propagation of each realization

$$\mathbf{Y}_{[j,i]} := \chi_i(t_{k+j}) = \underbrace{\mathcal{M}(\mathcal{M}(\dots \mathcal{M}(\chi_{true}(t_k) + \varepsilon_i(t_k))))}_{j \text{ times}}, \quad j = 1, \dots, N_F, i = 1, \dots, m, \quad (10)$$

Where N_F is the total number of forecast steps and $t_{k+1} - t_k = \Delta$. These trajectories can be sent directly to the stochastic optimizer. Alternatively, we can compute the multi-step forecast mean \mathbf{Y}^P and covariance matrix \mathbf{V}^P using the ensemble approximations (5). Note that the ensemble approach is able to capture the strong model nonlinearities more accurately. To obtain the ensembles, however, we need to specify the posterior covariance $\mathbf{V}(t_k)$. Because of the extremely large state dimensionality, the posterior matrix cannot be computed using traditional methods such as the Kalman filter propagation. We next introduce a method to obtain an approximation of this matrix.

Remark: In weather modeling the most successful state estimation approaches have been EnKF and 4D-Var. In both of these approaches, the Gaussian assumption of the posterior and measurement and model errors is required to define the maximum likelihood function [30; 26; 36; 20]. This is justifiable by the additive nature of many of error components and the central limit theorem as well as the convenience of normal distributions [30]. With this, the distribution of the current estimated state (posterior) is also implicitly assumed to be Gaussian. However, the forecast ensembles do not need to be Gaussian since the samples of the posterior are propagated through the nonlinear WRF model.

3.1.2 The NCEP Method for Covariance Estimation

The National Centers for Environmental Prediction (NCEP) method [43; 23; 31] has been used to estimate the *spatial* uncertainty information and estimate the posterior distribution $\mathcal{N}(0, \mathbf{V}(t_k))$. The idea is to estimate characteristic correlation distances to construct an empirical covariance matrix. The inferred characteristic horizontal correlation distance for this case is approximated by $L_H = 2$ degrees and by $L_V = 500$ meters in the vertical direction. The spatial correlation function between two spatial points $\chi(t_k, x_i, y_i, z_i)$ and $\chi(t_k, x_j, y_j, z_j)$ is defined as

$$C_{i,j} = \exp \left(-\frac{(x_j - x_i)^2 + (y_j - y_i)^2}{L_H^2} - \frac{(z_j - z_i)^2}{L_V^2} \right). \quad (11)$$

The correlation function (11) is used to construct the empirical covariance matrix from which the ensemble for the initial state field is drawn. Here, we focus on the temperature field $T(t_k, x, y, z)$. The true initial temperature field is not known exactly, but we assume that it is correctly represented by an unbiased random vector $\varepsilon_T(t_k)$. With this, the temperature field $T(t_k)$ is characterized by a random vector $T^B(t_k)$ with the following properties:

$$\begin{aligned} T^B(t_k) &= T(t_k) + \varepsilon_T(t_k), \quad \mathbf{E} [\varepsilon_T(t_k)\varepsilon_T(t_k)^T] = \mathbf{V}_{TT}(t_k), \quad \mathbf{E} [\varepsilon_T(t_k)] = 0 \\ &\Rightarrow T^B(t_k) \in \mathcal{N}(T(t_k), \mathbf{V}_{TT}(t_k)). \end{aligned}$$

The initial temperature field is approximated by an m -member ensemble drawn from

$$T_j^B(t_k) = T(t_k) + GCG^T\xi_j, \quad 1 \leq j \leq m, \quad \xi \in \mathcal{N}(0, 1),$$

where $GCG^T \approx \mathbf{V}_{TT}(t_k)$. Here, matrix G transforms the unbalanced variables into full quantities for temperature and is defined as $G_{[j]} := \sigma_G(z_j)\mathbb{I}$, $j = 1, \dots, nz$, and $G = \text{diag}(G_{[1]}, \dots, G_{[nz]})$, where

$$\sigma_G(z_j) = \mathbf{E} [T(t_k, x, y, z_j)] / \max_i (\mathbf{E} [T(t_k, x, y, z_i)]).$$

This covariance can then be used to compute the disturbance realizations (10).

3.2 Gaussian Process Modeling

A straightforward disturbance forecast alternative is to use historical measurement data to construct regression models. Consequently, an important question is if it is worth considering a highly sophisticated weather model to obtain forecast information. In this section, we present a regression modeling technique in order to establish a basis for comparison.

Several empirical modeling techniques can be used to generate weather forecast trends. An approach that has recently received attention is Gaussian process (GP) modeling [44; 48; 32; 42]. 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 [5]. We have found that this feature makes the GP approach more flexible. Consequently, this is the approach considered in this work. To illustrate the use of this technique, we construct a forecast model for the ambient temperature by regressing the future temperature (output) χ_{k+1} to the current and previous temperature values (inputs) $\chi_k, \dots, \chi_{k-N}$ that can be obtained from weather information data bases. In this case, N is selected long enough to capture the periodic trends of the ambient temperature. We define the model inputs as $\mathbf{X}_{[j]} = [\chi_{k-N-j}, \dots, \chi_{k-j}]$ and the outputs as $\mathbf{Y}_{[j]} = \chi_{k+1-j}$, and we collect a number of training sets $j = 0, \dots, N_{train}$. We assume that the inputs are correlated through an exponential covariance function of the form

$$\mathbf{V}(\mathbf{X}_{[j]}, \mathbf{X}_{[i]}, \eta) := \eta_0 + \eta_1 \cdot \exp\left(-\frac{1}{\eta_2} \|\mathbf{X}_{[j]} - \mathbf{X}_{[i]}\|^2\right), \quad i = 0, \dots, N_{train}, j = 0, \dots, N_{train}, \quad (12)$$

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

$$\log p(\mathbf{Y}|\eta) = -\frac{1}{2} \mathbf{Y}^T \mathbf{V}^{-1}(\mathbf{X}, \mathbf{X}, \eta) \mathbf{Y} - \frac{1}{2} \log \det(\mathbf{V}(\mathbf{X}, \mathbf{X}, \eta)). \quad (13)$$

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 temperature trends. The resulting GP posterior distribution is

$$\mathbf{Y}^P = \mathbf{V}(\mathbf{X}^P, \mathbf{X}, \eta^*) \mathbf{V}^{-1}(\mathbf{X}, \mathbf{X}, \eta^*) \mathbf{Y} \quad (14a)$$

$$\mathbf{V}^P = \mathbf{V}(\mathbf{X}^P, \mathbf{X}^P, \eta^*) - \mathbf{V}(\mathbf{X}^P, \mathbf{X}, \eta^*) \mathbf{V}^{-1}(\mathbf{X}, \mathbf{X}, \eta^*) \mathbf{V}(\mathbf{X}, \mathbf{X}^P, \eta^*). \quad (14b)$$

The inverse of the input covariance $\mathbf{V}_{\mathbf{X}} := \mathbf{V}^{-1}(\mathbf{X}, \mathbf{X}, \eta^*)$ (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

$$\mathbf{Y}^P = \mathbf{GP}(\mathbf{X}^P, \eta^*, \mathbf{V}_{\mathbf{X}}). \quad (15)$$

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 $j = 1, \dots, N_F$ do,

(a) Set $\mathbf{X}_{[j]}^P \leftarrow [\chi_{k-N}, \chi_{k-N+1}, \dots, \chi_k]$

(b) Compute $\mathbf{Y}_{[j]}^P = \mathbf{GP}(\mathbf{X}_{[j]}^P, \eta^*, \mathbf{V}_{\mathbf{X}})$

(c) Drop last measurement, set $\chi_{k+1} \leftarrow \mathbf{Y}_{[j]}^P$, 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 (14b).

This recursion generates the forecast mean $\mathbf{Y}^P = [\bar{\chi}_{k+1}, \dots, \bar{\chi}_{k+N_F}]$ and associated covariance matrix \mathbf{V}^P . Note that this disturbance trends are local (single point in space), as opposed to those obtained with the weather prediction model.

3.3 Validation of Uncertainty Models

We next validate the forecast information obtained from the NWP and 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 in this study. The data were obtained from the National Weather Service Office [38]. The temperature trends are presented in Figure 2. Note that strong temperature variations arise at different time scales (daily and seasonal). The variability is particularly strong during the winter.

To illustrate the forecasting capabilities of the GP modeling technique, we used a total of 120 training sets and we set $N = 24$. We consider a single-step strategy $N_F = 1$ and a multi-step strategy with $N_F = 20$. In Figure 3, we present the forecast mean and 100 samples drawn from the corresponding normal distributions $\mathcal{N}(\mathbf{Y}^P, \mathbf{V}^P)$. The forecast window corresponds to Aug. 1–6, 2006. In the top graph, we can see that the single-step strategy provides reasonable forecasts and the uncertainty bounds encapsulate the true temperature realizations. In the bottom graph, we can see that the multi-step GP model is able to capture the periodicity of the trends. However, the mean drifts away from the true temperature realizations and, more importantly, *the uncertainty bounds are not able to encapsulate the actual realizations*. This limits the application of this approach from an stochastic optimization point of view. Note that the ambient temperature follows strong variations as a result of spatial interactions and long-term meteorological phenomena that cannot be taken into account through empirical modeling techniques. Nevertheless, we emphasize that GP is quite effective for short-term forecasts. This suggests that the empirical modeling techniques could be a valuable for high-frequency control (on the order of seconds, minutes). Similar observations have been made in the context of short-term wind and solar radiation forecasting [24; 54].




Figure 2: Ambient temperature in Pittsburgh PA, 2006.

We next validate the forecast and uncertainty information of the NWP model. In the left graph of Figure 4, we show a multi-step ensemble of 30 ensemble members for temperature realizations, the expected temperature value, and measurements for five days (Aug. 1–6, 2006). This corresponds to a total of $N_F = 120$ forecast steps with $\Delta = 1\text{hr}$. In the right graph we present the reconstructed forecast distribution (mean $\pm 3\sigma$) obtained from the ensembles. Note that the forecast errors are small ($\pm 5^\circ\text{C}$) and the uncertainty




Figure 3: Temperature forecasts with single-step (top) and multi-step GP model (bottom). Forecast mean is solid line, samples are in light gray, and markers are actual realizations.

envelope encloses the true (measured) solution. In addition, the model can capture long-term temperature trends except at the third day, where an unusual temperature drop is observed. In Figure 5(a), we present the hourly evolution of the standard scores of the actual realizations for the multi-step GP model and for the weather model. In Figure 5(b), we present the cumulative standard error of the actual realizations for the multi-step GP model, for the weather model and for the standard normal. For example, a standard error less than 2 is obtained 65% of the time by the weather model but only 25% by the GP model. We conclude that the NWP uncertainty model is far more accurate and consistent.

The ensemble forecast also provides information of the spatial forecast error distribution. In Figure 6, we illustrate the horizontal correlation field for the temperature error in the Pittsburgh area corresponding to 10 a.m. August 1 and August 2, 2006 [38]. Note that the error field widens in time as the uncertainty of the forecast increases. Note also that strong temperature variations can arise in relatively narrow regions. Therefore, we emphasize that accounting for spatial effects is critical for accurate and consistent forecasts.




Figure 4: 5-day-ahead temperature forecast of NWP for the Pittsburgh airport area during August 1–6, 2006. Forecast mean and 30 ensembles obtained from perturbations of initial temperature field (left). Reconstruction of forecast distribution from ensembles - mean $\pm 3\sigma$ (right).

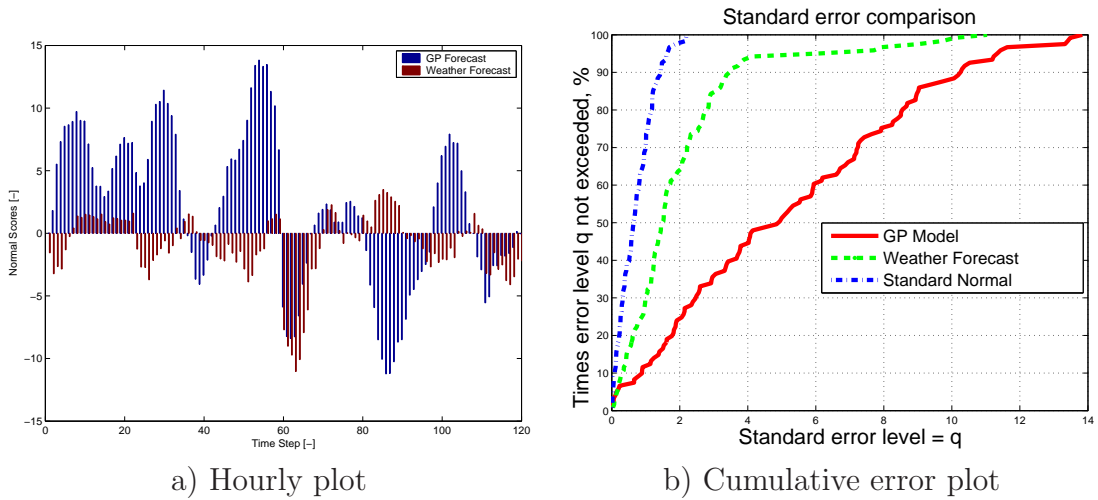


Figure 5: Standard scores for GP and weather model forecasts.

4 Integrative Study for Building System

In this section, we present a D-RTO simulation study on the climate control of a building system. Our objective is to illustrate how the use of forecast information can translate into lower operating costs. Finally, we analyze the performance of the proposed uncertainty quantification techniques when coupled to the stochastic optimization formulation.




Figure 6: Correlation field for the temperature errors in the Pittsburgh area at 10 a.m. on consecutive days.

4.1 Economic Impact of Forecasting

Commercial buildings are energy-intensive facilities where considerable cost savings can be realized through optimal operating strategies. As an example, researchers have found that the thermal mass of a building can be used for temporal energy storage [6]. With this, one can 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. For instance, a cooling strategy that has been used in commercial facilities consists in cooling down the building as much as possible at night when electricity is cheaper so as to reduce the amount of cooling needed during the day when electricity is more expensive [7]. 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, we point out that the optimal timing at which it is decided to start the cooling at night directly depends on the ambient temperature expected the next day. In addition, because of governmental regulations, special care needs to be taken to stay within the thermal comfort zone at all times. Consequently, incorporating weather forecast information in the optimization/control formulation can be beneficial. The building system under consideration is sketched in Figure 7. We assume a total volume of $10,000 \text{ m}^3$ and a total surface area of $3,500 \text{ m}^2$. The building is equipped with a gas furnace, an electric heater, and an electric cooling system. The ambient temperature information used in this study is presented in Figure 2. The dynamic response of the building internal temperature is modeled by an ordinary differential equation; the building wall is modeled by a second-order PDE that accounts for conductive effects along the wall. The ambient temperature enters the model through a Robin boundary condition at the wall external face. The basic heat-transfer model structure has been obtained from [10]. To analyze the effect of adding forecast information of the ambient temperature we first solve an open-loop dynamic optimization problem with perfect forecast information




Figure 7: Schematic representation of building integration with heating, ventilation, and air-conditioning (HVAC) system.

and a prediction horizon of one year. The optimization problem has the following form:

$$\min_{\varphi_c^{elec}(\tau), \varphi_h^{gas}(\tau), \varphi_h^{elec}(\tau)} \int_{t_k}^{t_k+T} [C_{elec}(\tau)\varphi_c^{elec}(\tau) + C_{elec}(\tau)\varphi_h^{elec}(\tau) + C_{gas}\varphi_h^{gas}(\tau)] d\tau$$

$$C_I \cdot \frac{\partial T_I}{\partial \tau} = \varphi_h^{gas}(\tau) + \varphi_h^{elec}(\tau) - \varphi_c^{elec}(\tau) - S \cdot \alpha' \cdot (T_I(\tau) - T_W(\tau, 0)) \quad (16a)$$

$$\frac{\partial T_W}{\partial \tau} = \beta \cdot \frac{\partial^2 T_W}{\partial x^2} \quad (16b)$$

$$0 = \alpha' (T_I(\tau) - T_W(\tau, 0)) + \mathbf{k} \cdot \frac{\partial T_W}{\partial x} \Big|_{(\tau, 0)} \quad (16c)$$

$$0 = \alpha'' (T_W(\tau, L) - T_A(\tau)) + \mathbf{k} \cdot \frac{\partial T_W}{\partial x} \Big|_{(\tau, L)} \quad (16d)$$

$$T_I^{min} \leq T_I(\tau) \leq T_I^{max} \quad (16e)$$

$$0 \leq \varphi_h^{gas}(\tau), \quad 0 \leq \varphi_h^{elec}(\tau), \quad 0 \leq \varphi_c^{elec}(\tau) \quad (16f)$$

$$T_I(0) = T_I^k \quad (16g)$$

$$T_W(0, x) = T_W^k(x), \quad (16h)$$

where $T_A(\tau)$ is the ambient temperature, $T_I(\tau)$ is the internal temperature, and $T_W(\tau, x)$ is the wall temperature (all of them in °C). The controls are the gas heating power $\varphi_h^{gas}(\tau)$, the electric heating power $\varphi_h^{elec}(\tau)$, and the electric cooling power $\varphi_c^{elec}(\tau)$ (all of them in kcal/hr). The model parameters are summarized in Table 1. The base wall thickness is assumed to be 0.20 m. We assume an on-peak electricity price of 0.12 \$/kWh available from 9 a.m. to 10 p.m. The off-peak price is 0.04 \$/kWh. A demand rate of 16 \$/kW is charged for the monthly peak electricity demand. The natural gas price is fixed at 0.10 \$/kWh. Average prices were obtained from [51]. The thermal comfort zone is assumed

to be 69-77°F. The above PDE-constrained optimization problem is discretized by using a central difference scheme in the axial dimension and an implicit Euler scheme in time. The resulting LP was implemented in AMPL and solved with IPOPT [52].

From the solution of the open-loop dynamic optimization problem, we obtain the optimal cost and use it as a reference for the best economic performance of the system. The resulting minimum annual cost is \$28,672 (demand cost is approximately 60% of total cost). The one-year forecast problem contains 96,613 constraints and 26,349 degrees of freedom and can be solved in 25 iterations and 30 CPU-seconds with IPOPT. All numerical calculations are performed on a personal computer with 4 GB of memory and a Duo-Core Intel processor running at 2.1 GHz. We then solve closed-loop D-RTO problems over the entire year with prediction horizons of 1, 3, 6, 9, 12, 16, and 24 hr. An update period Δ of 1 hr is used. The 24-hr forecast problem contains 253 constraints and 70 degrees of freedom and can be solved, in warm-start mode, in 10 iterations and 0.1 CPU-seconds. The relative costs (excluding demand costs) are presented in Figure 8(a). As can be seen, for a purely reactive strategy, the relative costs can go as high as 24% as a result of lack of proactiveness. In addition, we observe that a horizon of 24 hr is sufficient to achieve the minimum potential costs. The reason is that the thermal mass of the building cannot be used for a very long time because energy is lost through the wall. In fact, we found that as the building insulation is enhanced, the costs can be further reduced. To illustrate this situation, in Figure 8(b) we present the relative costs with an increased wall thickness of 0.3 m. As can be seen, using a forecast of 24 hr can reduce costs by 45%. 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*, which is in turn influenced by the building characteristics. The predicted cost savings agree with the results of a previous economic study on a photovoltaic-hydrogen hybrid system [54]. In that study, we found that the operating costs can be reduced by as much as 75% by incorporating forecast information of the solar radiation.

In Figure 9 we present the temperature set-points for the 24-hr and 1-hr forecast cases during 10 days in the winter season. As can be seen, the 24-hr forecast strategy determines the optimal timing at which electric heating needs to be turned on at night. Note that the optimum timing and the peak temperature depend on the expected ambient temperature. On the other hand, the reactive strategy is not able to foresee the structure of the electricity prices. This strategy suggests that the optimal policy is to keep the temperature set-point always at the lowest possible value in order to reduce the overall heating costs. Although this strategy seems intuitive, it is clearly not optimal if the structure of the electricity rates and the thermal mass of the building can be exploited. From Figure 10, we observe that the optimal cooling policy during the summer follows a peak-shifting strategy. The resulting policy recommends letting the building cool down at night until the temperature gets close to the lower limit of the comfort zone. During the day, the building is allowed to heat up progressively until it reaches the highest limit of the comfort zone. Similar results have been obtained by Braun and coworkers [7]. The proposed D-RTO framework can account for time variations and correct the policy

Table 1: Building model parameters.

Parameter	Value	Units	Meaning
β	0.001	$\frac{m^2}{hr}$	thermal diffusivity of wall
C_I	8,325	$\frac{kJ}{^\circ C}$	internal heat capacity
\mathbf{k}	1.16×10^{-4}	$\frac{kW}{m \cdot ^\circ C}$	conductivity of wall
S	3,500	m^2	wall total surface area
A	1,000	m^2	usable total surface area
V	10,000	m^3	building total volume
α'	4.64×10^{-3}	$\frac{kW}{m^2 \cdot ^\circ C}$	convective heat transfer coefficient (wall inner side)
α''	1.16×10^{-2}	$\frac{kW}{m^2 \cdot ^\circ C}$	convective heat transfer coefficient (wall outer side)
L	0.20	m	wall thickness
C_{elec}	0.12	$\frac{\$}{kWh}$	on-peak electricity cost
C_{elec}	0.04	$\frac{\$}{kWh}$	off-peak electricity cost
C_{gas}	0.10	$\frac{\$}{kWh}$	natural gas cost

automatically on-line. In this simplified study the cooling requirements are negligible because we account only for heat gains and losses through the wall. In addition, the day-night temperature difference at this location is large during summer, as seen in Figure 10. A more detailed study should also account for internal heat gains, radiation heating, air recycling, and humidity factors. Nevertheless, these preliminary results indicate that the performance of operating strategies can benefit from anticipating the weather conditions.




Figure 8: Impact of forecast horizon on economic performance of building system.




Figure 9: Optimal temperature set-points of closed-loop D-RTO with 1-hr and 24-hr forecasts. Comfort zone is highlighted in gray.




Figure 10: Internal temperature set-point and ambient temperature during 10 days in summer. Closed-loop D-RTO with a forecast of 24 hr was used. Comfort zone is highlighted in gray.

4.2 Stochastic Optimization Results

The previous study assumes that the weather information is perfectly known. We now consider the case in which the temperature trends are obtained from the forecast models. We solve the SAA approximation of the stochastic counterpart of problem (16) over a horizon of 5 days. The SAA problem is given in equation (17). We use 100 samples drawn from the forecast distributions of the multi-step GP model and of the weather model shown in Figures 3 and 4, respectively. After discretization, the resulting NLP contains 130,900 constraints and 357 degrees of freedom. The problem can be solved, in warm-start mode, in 20 iterations and 68 CPU-seconds with IPOPT. The resulting open-loop profiles for the building internal temperature are presented in Figure 11. In the top graph, we present the temperature profile for the ideal strategy where perfect forecast information is assumed. Since there is no uncertainty, the predicted temperature profile matches the actual realization. Note that the optimal set-point policy 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

temperature profiles obtained using forecast information from the GP model. The gray lines are the predicted realizations of the dynamic model in the SAA formulation. This is an outcome of the solution of the SAA problem. Note that, since the uncertainty structure provided by GP is not able to capture the ambient temperature, the actual realization of the internal building temperature goes outside the comfort zone. In the bottom graph, we see that the use of weather model forecast results in an temperature trajectory that stays within the comfort zone at all times. Note that the variance of the predicted temperature realizations increases with time. In addition, since the comfort zone is very narrow ($\approx 5^\circ\text{C}$), high-precision forecast information is needed to realize economic benefits.

The cost penalty sustained by the NCEP weather uncertainty approach from subsections 3.1 and 3.1.2 when compared to the GP model uncertainty approach is about 10-20% of the ideal cost. One should bear in mind, however, that the GP model cost turns out to be infeasible for the actual realizations (the building temperature significantly exits the comfort level), so using only cost as a performance is misleading in this case. We could easily imagine some financial measure of the violation and report it to balance the perceived cost drop. Nevertheless, given the complex regulatory nature of the comfort level limits, their violation cost may easily be understated. For example, the 10CFR434 federal regulations for new federal, commercial and multi-family high-rise residential buildings in the United States require compliance with the comfort zone at least 98 % of the time the building is occupied. Therefore using the GP model would result in the operator being in violation of the federal law (the constraint violation in Figure 11 would be out of compliance more than 30 % of the period stated), the cost of which is difficult to fully assess. Given the difficulty of pricing the violation, it is more beneficial to regard the situation from the constrained optimization perspective and state that feasibility takes precedence over low cost. We conclude that the weather uncertainty model is the only one that has a sufficiently accurate description of the uncertainty to result in a feasible policy at a cost that is still substantially lower than the reactive policy cost. Finally, we should emphasize that the above results neglect the presence of a back-up controller that could potentially bring back the temperature within the thermal comfort zone. While this is certainly the case in practice, an economic penalty will have to be considered for the back-up controller as well. This would require a more detailed closed-loop case study.

$$\min_{\varphi_c^{elec}(\tau), \varphi_h^{gas}(\tau), \varphi_h^{elec}(\tau)} \frac{1}{N_S} \sum_{j=1}^{N_S} \int_{t_k}^{t_k+T} [C_{elec}(\tau)\varphi_c^{elec}(\tau) + C_{elec}(\tau)\varphi_h^{elec}(\tau) + C_{gas}\varphi_h^{gas}(\tau)] d\tau$$

$$C_I \cdot \frac{\partial T_I^j}{\partial \tau} = \varphi_h^{gas}(\tau) + \varphi_h^{elec}(\tau) - \varphi_c^{elec}(\tau) - S \cdot \alpha' \cdot (T_I^j(\tau) - T_W^j(\tau, 0)) \quad (17a)$$

$$\frac{\partial T_W^j}{\partial \tau} = \beta \cdot \frac{\partial^2 T_W^j}{\partial x^2} \quad (17b)$$

$$0 = \alpha' (T_I^j(\tau) - T_W^j(\tau, 0)) + \mathbf{k} \cdot \frac{\partial T_W^j}{\partial x} \Big|_{(\tau, 0)} \quad (17c)$$

$$0 = \alpha'' (T_W^j(\tau, L) - T_A^j(\tau)) + \mathbf{k} \cdot \frac{\partial T_W^j}{\partial x} \Big|_{(\tau, L)} \quad (17d)$$

$$T_I^{min} \leq T_I^j(\tau) \leq T_I^{max} \quad (17e)$$

$$0 \leq \varphi_h^{gas}(\tau), \quad 0 \leq \varphi_h^{elec}(\tau), \quad 0 \leq \varphi_c^{elec}(\tau) \quad (17f)$$

$$T_I(0) = T_I^k \quad (17g)$$

$$T_W(0, x) = T_W^k(x), \tau \in [t_k, t_k + T], j = 1, \dots, N_S. \quad (17h)$$

5 Conclusions and Future Work

In this work, we demonstrate that significant costs reductions can be achieved by using on-line optimization strategies that can anticipate the weather conditions. In particular, we show that adding forecast information provides a mechanism to compute proactive operating policies that can lead to enhanced performance.

We present different strategies to obtain weather forecast information. We emphasize that empirical models provide quick estimates of the weather trends but they are limited to short horizons and can lead to inconsistent uncertainty bounds. Motivated by these facts, we discuss the potential of using detailed weather models to obtain forecasts. We demonstrate that these models are capable of providing more accurate forecasts and are able to capture temporal and spatial correlations of the state fields. We extend a weather model to provide forecast covariance information through the ensemble approach.

As future work, we are interested in establishing a full connection between the weather model forecasts and the stochastic dynamic optimization framework. To do so, we first must implement the ensemble-based approach in a closed-loop manner. Since the weather model is extremely computationally expensive, a dedicated, centralized parallel computing architecture is needed. In addition, since the amount of data to be handled is huge, strategies must be established to communicate only the essential forecast statistical information. Another important issue is the fact that the weather model provides information over relatively coarse fields that need to be mapped to the specific location of the system under consideration. To this end, we are interested in using a Gaussian process modeling




Figure 11: Performance of weather forecast-based operating strategies. Thermal comfort zone is highlighted by thick solid lines, predicted temperatures are gray lines, and actual realizations are dashed lines.

framework to interpolate the spatiotemporal fields. We are also interested in addressing the complexity of large-scale stochastic programming problems through adaptive sampling and variance reduction techniques. Finally, establishing potential economic benefits in different applications such as polygeneration systems is an important research area.

Acknowledgments

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

References

- [1] M. Anitescu. Spectral stochastic finite element methods for parametric constrained optimization problems. *SIAM Journal of Numerical Analysis*, 47:1739–1759, 2009.
- [2] M. Anitescu and J. Birge. Convergence of stochastic average approximation for stochastic optimization problems with mixed expectation and per-scenario constraints. *Submitted for Publication*, 2009.

- [3] J. Arlington, C. C. Pantelides, B. Rustem, and B. A. Tanyi. An algorithm for constrained nonlinear optimization under uncertainty. *Automatica*, 35:217–228, 1999.
- [4] D. Betsekas and C. M. White. Dynamic programming and stochastic control. *IEEE Transactions on Systems, Man and Cybernetics*, 7:758–759, 1977.
- [5] G. E. P. Box, G. Jenkins, and G. Reinsel. *Time Series Analysis: Forecasting and Control*. Prentice-Hall, New Jersey, 1994.
- [6] J. E. Braun. Reducing energy costs and peak electricity demand through optimal control of building thermal storage. *ASHRAE Transactions*, 96(2):876–888, 1990.
- [7] J. E. Braun, K. W. Montgomery, and N. Chaturvedi. Evaluating the performance of building thermal mass control strategies. *HVAC&Research*, 7(4):403–428, 2001.
- [8] G. Burgers, P.J. van Leeuwen, and G. Evensen. Analysis scheme in the ensemble Kalman filter. *Monthly Weather Review*, 126:1719–1724, 1998.
- [9] J. Busch, J. Oldenburg, M. Santos, A. Cruse, and W. Marquardt. Dynamic predictive scheduling of operational strategies for continuous processes using mixed-logic dynamic optimization. *Comp. Chem. Eng.*, 31:574–587, 2007.
- [10] C. Carrara, L. F. Fantini, and A. Lorenzi. Influence of the outdoor temperature variation on the performance of a building heating system. In C. J. Hoogendoorn, editor, *Energy Conservation in Heating, Cooling and Ventilating Buildings*. 1978.
- [11] A. Collazos, F. Marchal, and C. Gaehler. Predictive optimal management method for the control of polygeneration systems. *Computers and Chemical Engineering, In Press*, 2009.
- [12] E.M. Constantinescu, T. Chai, A. Sandu, and G.R. Carmichael. Autoregressive models of background errors for chemical data assimilation. *Journal of Geophysical Research*, 112:D12309, 2007.
- [13] E.M. Constantinescu, A. Sandu, T. Chai, and G.R. Carmichael. Ensemble-based chemical data assimilation I: General approach. *Quarterly Journal of Royal Meteorological Society*, 133(626):1229–1243, July 2007.
- [14] E.M. Constantinescu, A. Sandu, T. Chai, and G.R. Carmichael. Ensemble-based chemical data assimilation II: Covariance localization. *Quarterly Journal of Royal Meteorological Society*, 133(626):1245–1256, July 2007.
- [15] P. Courtier, E. Andersson, W. Heckley, J. Pailleux, D. Vasiljevic, M. Hamrud, A. Hollingsworth, F. Rabier, and M. Fisher. The ECMWF implementation of three-dimensional variational assimilation (3D-Var). I:Formulation. *Quarterly Journal of the Royal Meteorological Society*, 124(550):1783–1807, 1998.

- [16] P. Courtier, J.-N. Thepaut, and A. Hollingsworth. A strategy for operational implementation of 4D-Var, using an incremental approach. *Quarterly Journal of the Royal Meteorological Society*, 120:1367–1387, 1994.
- [17] M. Diehl, H. G. Bock, and E. Kostina. An algorithm for constrained nonlinear optimization under uncertainty. *Mathematical Programming*, 107:213–230, 2006.
- [18] J.L. Doob. *Stochastic processes*. 1953.
- [19] A. P. Douglas, A. M. Breipohl, F. N. Lee, and R. Adapa. Risk due to load forecast uncertainty in short term power system planning. *IEEE Transactions on Power Systems*, 13:1493–1499, 1998.
- [20] AWF Edwards. *Likelihood*. Johns Hopkins University Press, 1992.
- [21] G. Evensen. Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research*, 99(C5):10143–10162, 1994.
- [22] M. Fisher, J. Nocedal, Y. Tremolet, and S. J. Wright. Data assimilation in weather forecasting: A case study in PDE-constrained optimization. *Submitted for Publication*, 2008.
- [23] R. Franke and E. Barker. Vertical correlation functions for temperature and relative humidity errors. *Monthly Weather Review*, 128(12):3962–3981, 2000.
- [24] G. Giebel. Development of a next generation wind resource forecasting system for the large-scale integration of onshore and offshore wind farms. *Technical Report Project ANEMOS*, 2002.
- [25] N. Groewe-Kuska, M. P. Nowak, and I. Wegner. Modeling of uncertainty for real-time management of power systems. In *Online Optimization of Large-Scale Systems*, pages 623–648. Springer, Berlin, 2001.
- [26] PL Houtekamer, L. Lefaiivre, J. Derome, H. Ritchie, and H.L. Mitchell. A system simulation approach to ensemble prediction. *Monthly Weather Review*, 124(6):1225–1242, 1996.
- [27] P.L. Houtekamer and H.L. Mitchell. Data assimilation using an ensemble Kalman filter technique. *Monthly Weather Review*, 126(3):796–811, 1998.
- [28] R. Huang, S. C. Patwardhan, and L. T. Biegler. Multi-scenario-based robust nonlinear model predictive control with first principle models. In R. M. de Brito Alves, C. A. Oller do Nascimento, and E. C. Biscaia Jr., editors, *Proceedings of 10th International Symposium on Process Systems Engineering*, 2009.

- [29] J. Kadam, W. Marquardt, M. Schlegel, T. Backx, O. Bosgra, and P.J. Brouwer. Towards integrated dynamic real-time optimization and control of industrial processes. *Proceedings Foundations of Computer-Aided Process Operations (FOCAPO2003)*, pages 593–596, 2003.
- [30] E. Kalnay. *Atmospheric Modeling, Data Assimilation and Predictability*. Cambridge University Press, 2002.
- [31] 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(3):437–471, 1996.
- [32] Herbert KH Lee, Dave M Higdon, Catherine A Calder, and Christopher H Holloman. Efficient models for correlated data via convolutions of intrinsic processes. *Statistical Modeling*, 5(1):53–74, 2005.
- [33] J. Linderoth, A. Shapiro, and S. Wright. The empirical behavior of sampling methods for stochastic programming. *Annals of Operations Research*, 142(1):215–241, 2006.
- [34] L. Magni and R. Scattolini. Robustness and robust design of mpc for nonlinear systems. In *Assessment and Future Directions of NMPC*, pages 239–254. Springer, Berlin, 2007.
- [35] T.E. Marlin and A.N. Hrymak. Real-time operations optimization of continuous processes. In *Fifth International Conference on Chemical Process Control (CPC-5)*, 1996.
- [36] H.L. Mitchell and P.L. Houtekamer. Ensemble size, balance, and model-error representation in an ensemble Kalman filter. *Monthly Weather Review*, 130:2791–2808, 2002.
- [37] Z. Nagy and R. Braatz. Distributional uncertainty analysis using power series and polynomial chaos expansions. *Journal of Process Control*, 17:229–240, 2007.
- [38] National Oceanic and Atmospheric Association. National Weather Service Forecast Office, 2009. <http://www.erh.noaa.gov/pbz/hourlyclimate.htm>.
- [39] A. Nemirovski and A. Shapiro. Convex approximations of chance constrained programs. *SIAM J. Optim.*, 17:969–996, 2007.
- [40] F. Oldewurtel, C. N. Jones, and M. Morari. A tractable approximation of chance constrained stochastic mpc based on affine disturbance feedback. In *Proceedings of IEEE Conference on Decision and Control*, pages 4731–4736, 2008.
- [41] E. Ott, B.R. Hunt, I. Szunyogh, A.V. Zimin, E.J. Kostelich, M. Corazza, E. Kalnay, D.J. Patil, and J.A. Yorke. A local ensemble Kalman filter for atmospheric data assimilation. *Tellus A*, 56(5):415–428, 2004.

- [42] C.J. Paciorek and M.J. Schervish. Nonstationary covariance functions for Gaussian process regression. In *Advances in Neural Information Processing Systems 16: Proceedings of the 2003 Conference*, page 273. Bradford Book, 2004.
- [43] D.F. Parrish and J.C. Derber. The National Meteorological Center’s spectral statistical-interpolation analysis system. *Monthly Weather Review*, 120(8):1747–1763, 1992.
- [44] C. E. Rasmussen and C. K. Williams. *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, 2006.
- [45] J. B. Rawlings and R. Amrit. Optimizing process economic performance using model predictive control. In *Proceedings of the International Workshop on Assessment and Future Directions of NMPC*, 2008.
- [46] L.P. Riishojgaard. A direct way of specifying flow-dependent background error correlations for meteorological analysis systems. *Tellus A*, 50(1):42–42, 1998.
- [47] A. Ruszczyński and A. Shapiro. *Stochastic Programming (Handbooks in Operations Research and Management Series)*. Elsevier Science BV, Amsterdam, 2003.
- [48] Sujit K. Sahu, Stan Yip, and David M. Holland. Improved space-time forecasting of next day ozone concentrations in the eastern US. *Atmospheric Environment*, 43(3):494 – 501, 2009.
- [49] W.C. Skamarock, J.B. Klemp, J. Dudhia, D.O. Gill, D.M. Barker, W. Wang, and J.G. Powers. A description of the Advanced Research WRF version 2. Technical Report Tech Notes-468+ STR, NCAR, 2005.
- [50] O. Ulleberg. *Stand-Alone Power Systems for the Future: Optimal Design, Operation and Control of Solar-Hydrogen Energy Systems*. PhD thesis, Trondheim, Norway, 1998.
- [51] Kansas State University. Comparing fuel costs of heating and cooling systems. *Technical Report*, 2003.
- [52] A. Wächter and L. T. Biegler. On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming. *Math. Program.*, 106:25–57, 2006.
- [53] J. Wang, M. Shahidehpour, and Z. Li. Security-constrained unit commitment with volatile wind power generation. *IEEE Transactions on Power Systems*, 23:1319–1327, 2008.
- [54] V. M. Zavala, M. Anitescu, and T. Krause. On the Optimal On-Line Management of Photovoltaic-Hydrogen Hybrid Energy Systems. In R. M. de Brito Alves, C. A. Oller do Nascimento, and E. C. Biscaia Jr., editors, *Proceedings of 10th International Symposium on Process Systems Engineering*, 2009.

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.