Decision-Support Tools for Renewables-Rich Power Systems: A Stochastic Futures Approach*

Jiayi Jiang, Sandip Roy, Juhua Liu, and Vaibhav Donde**

Abstract. The growing penetration of intermittent renewables (primarily wind and solar generation) in deregulated electric power systems is introducing significant challenges in forecasting generation and scheduling units. At the same time, the pervasive integration of cyber- tools in the control room provides unique opportunities for leveraging data sources like weather forecasts, computational resources, and visualization tools for real-time decision-making. Here, we introduce a framework and algorithm set for day-ahead generation scheduling, or unit commitment, that takes advantage of the close tie between cyber- and physical- resources in the electric power grid. First, we use a class of stochastic automata models known as influence models to forecast relevant spatio-temporal environmental parameters (wind speeds/direction, cloud cover), and in turn simulate probabilistic wind and solar generation futures across a wide area. These models can be parameterized in real time to statistically match publicly-available ensemble forecast products, yet can be tailored to provide generation futures at appropriate spatial and temporal resolutions for scheduling. The models also permit rapid selection of representative renewable-generation futures, and are able to capture local variability and spatial/temporal correlation in the generation profiles. Second, a new method for unit scheduling for the day-ahead market, which uses the probabilistic wind/solar generation futures, is proposed and

---

Jiayi Jiang · Sandip Roy
Washington State University

Juhua Liu
ABB US Corporate Research

Vaibhav Donde
Pacific Gas & Electric

* This research described here was primarily supported by the ABB Research Grant Program.
** Work performed while at ABB US Corporate Research.

S.K. Khaitan, J.D. McCalley, and C.-C. Liu (eds.), *Cyber Physical Systems Approach to Smart Electric Power Grid,* Power Systems, DOI: 10.1007/978-3-662-45928-7_6

developed in a preliminary way. A novelty in this approach is a pre-selection step that can provide operators with situational awareness of critical (sensitive) units. The generation-scheduling and unit-commitment tools are demonstrated on a small-scale example, which is concerned with wind generation in the Columbia River Gorge of Washington State on a historical weather day.

1 Introduction

Electric-power-system operation requires coordinated scheduling and dispatch of generation units across a wide area, to match generation with demand. In many modern deregulated systems [1-13], scheduling and dispatch are achieved at three different time horizons. First, the on/off schedules and tentative hourly dispatch levels of generators are set by the transmission system operator (TSO) or independent system operator (ISO), usually via a binding market, on the day ahead. Second, refined dispatch levels are obtained via an hourly market mechanism which uses an economic dispatch. Finally, local small-scale mismatches are corrected for at a fast time scale, usually on the order 5-10 minutes.

The research described here is primarily concerned with generation scheduling for the day-ahead market. Historically, unit scheduling (as well as longer-term generation-resource planning) was done by human operators, who largely drew on experience to develop on/off schedules for a limited number of generator units. As electric power networks have become increasingly complex and computing technologies have improved, automation for *unit commitment* have been developed and integrated into transmission system operation. These unit commitment technologies, used in tandem with experience-driven decision-making, have proved valuable for wide-area management in both deregulated and regulated systems.

The last ten years has seen a rapid integration of intermittent renewable generation (primarily, wind and solar generation) into electric power systems worldwide, and the penetration of these intermittent renewables is expected to continue growing rapidly. These new generation technologies hold promise to permit sustainable low-cost power for years to come. However, they also bring forth new challenges in control and management of the power grid across multiple temporal and spatial scales, including specifically for day-ahead unit scheduling. Crucially, intermittent renewable generation trajectories are dependent on environmental parameters (e.g., wind speed and direction, cloud cover, humidity, etc.) which may have significant uncertainty at a 24-36 hour look-ahead horizon. This uncertainty must be accounted for in commitment and dispatch of conventional generation, and hence the unit-commitment problem becomes a stochastic one. In addition, the intermittence and consequent temporal variability in wind and solar generation means that unit schedules may change significantly from day to day. This variability makes experience-driven decision-making more difficult, and also requires flexible scheduling paradigms and improved tools for evaluating system-level performance (including economic performance, security and fault management, etc). As the penetration of intermittent renewables increases, these challenges in day-ahead scheduling will become increasingly prominent.

New tools for scheduling generation for the day-ahead market are needed to meet these challenges. These include tools for 1) forecasting intermittent-renewable generation futures, 2) stochastic unit commitment, and 3) evaluation of power-network performance across renewable-generation futures. Additionally, advances in these directions must be translated into practical decision-support software for the control room. In fact, numerous research efforts are underway in these directions. However, these efforts are still largely academic in nature, and have not yet been translated to implemented software solutions. Our viewpoint is that several barriers remain in obtaining implementable technologies:

1) Forecasts of uncertain environmental futures and consequent generation trajectories are needed, that have sufficient resolution for decision-making yet capture uncertain propagation across a wide area as needed for unit commitment.
2) Techniques for stochastic unit commitment are needed that yield practical, robust, and economically viable schedules across generation futures, yet are computationally attractive for wide-area scheduling.
3) End-to-end solutions are needed, that use realistic environmental forecasts for unit commitment and system performance evaluation.

While the growing penetration of intermittent renewables is complicating generation scheduling, new technologies also provide entirely new capabilities for resource scheduling that have not yet been fully exploited. During the last 20 years or so, a wide array of new computing and communication tools have been introduced in the control room: these include increasingly-powerful computers and sophisticated software for analysis, dedicated communication channels as well as high-speed Internet access, and mobile handheld technologies (cell phones, iPads, etc.), among others. These pervasive cyber- tools can facilitate control and management of the wide-area network across time scales [41,42]. In particular, relevant to the unit commitment problem for the renewables-rich grid, these technologies can allow fast transfer of high-dimensional weather-forecast data to the control room, provide operators with convenient interfaces and displays to evaluate consequences of decisions, simplify wide-area monitoring, permit rapid integration of stakeholders' inputs, and allow intensive computing for weather-impact forecasting and schedule optimization. Indeed, the new cyber-technologies have brought about rapid advances in control room operations, but they have not yet yielded significant improvements in unit commitment for a renewables-rich grid. At its essence, exploiting these technologies for stochastic unit commitment requires an understanding of the tight interface between engineered (electromechanical), natural-world (weather), human (market and operational), and cyber components in the electric power grid.

The research presented here approaches stochastic unit commitment from this "cyber-physical systems" viewpoint, focusing particularly on cyber- solutions to the forecasting and scheduling aspects of the problem. Research efforts in three directions are discussed:

1) Development of an end-to-end operational concept for day-ahead unit scheduling.
2) Motivation for and development of a new generation-forecasting tool, which uses a stochastic automaton model known as the influence model.
3) Exploration of tools for stochastic unit commitment that use the new generation-forecasting model.

The chapter is organized as follows. The end-to-end operational concept is first introduced (Section 2). Next, the new generation-forecasting tool is developed in detail (Section 3), and illustrated using a case study of wind generation in the Columbia River Gorge area of Washington State on a historical weather day. Finally, some initial explorations on using the generation forecasts for stochastic unit commitment are presented (Section 4), and conclusions are given (Section 5).

2 Operational Concept

A *stochastic-futures approach* to generation scheduling for the day-ahead market is considered, see Figure 1. The operational concept has two main parts: a module for determining representative spatiotemporal futures (or time-trajectories of wind and solar generation over a 48-hour horizon (the *generation-forecasting module,* as contained in the green blocks), and a second module for scheduling dispatchable generation units using these wind/solar generation futures (the *stochastic unit commitment module*, pink blocks). These modules require development of new algorithms for generation forecasting and unit commitment, respectively, as well as prototype Matlab software development.

The generation-forecasting module in our approach exploits a new *influence modeling* technology (14-17). This technology leverages ensemble forecast outputs, but also allows interpolation of forecasts to the proper resolution for generation forecasting, represents spatial and temporal correlation in weather/generation, and permits rapid simulation of many generation futures. The module uses the influence-modeling technology as follows (see flowchart): 1) relevant forecast data (wind speeds and directions, cloud cover, humidity) is extracted from an ensemble forecast (e.g., the Short Range Ensemble Forecast or SREF, which is available in the public domain, see 18); 2) influence models for wind speeds and cloud dynamics are built (parameterized) to statistically match the ensemble-forecast data; 3) many possible spatio-temporal futures of wind- and solar- generation are obtained through simulation of the stochastic influence model and mapping of the results into generation profiles; and 4) a few representative futures are chosen using the probabilistic-collocation method

(19,20). The futures produced by the module predict wind/solar generation at each network bus over the full day ahead, at a 15-min. resolution.

Meanwhile, the scheduling module aims to develop day-ahead commitment plans for dispatchable (non-renewable) generation, to minimize an expected performance cost across the representative scenario set while respecting numerous constraints (including ramp-up and ramp-down constraints, and transmission-network constraints upon dispatch). The performance cost for scheduling in our formulation captures dispatch cost, ramp-up and ramp-down costs, reserve-generation usage costs, and line losses.

Since the forward market for the day ahead is a binding one, the module must either provide a single plan or a very small number of alternatives (with human operators choosing one). To combat the computational challenges inherent to stochastic unit commitment (UC) and to provide operators with robust plans, a new two-stage approach to generator scheduling is proposed, which contrasts with existing stochastic UC paradigms (11,21,22). In particular, rather than trying to optimize all generators' schedules (including when they are on-line/off-line and hourly dispatch levels) at once, we instead first **identify critical units** that are difficult to plan. These are units whose optimal schedules of on-line times and/or dispatch are highly sensitive to the future renewable-generation profile. The critical units are identified by pursuing hourly economic dispatch (EDs) for each representative weather-impact future assuming all units are potentially online, and determining the units whose generation levels are sensitive to the weather future. In the second stage, the full scheduling problem is

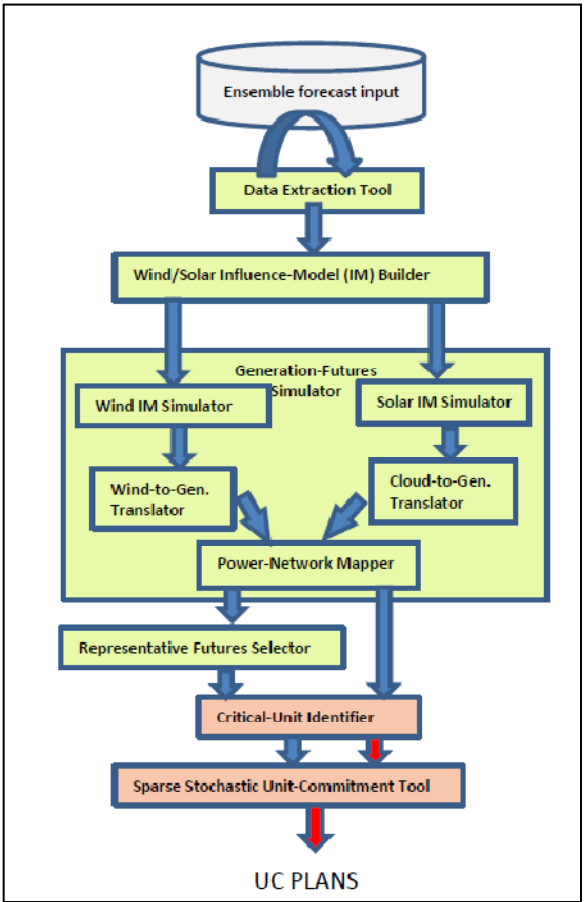

Fig. 1 Operational concept for the stochastic-futures-based unit commitment tool

solved using only the critical units' on/off times and dispatch as design variables, while using mean dispatch levels from the first stage for the remaining non-critical units.

Several new algorithms are needed for the stochastic-futures-based UC solution, including for building and simulating the influence model, translating weather futures into renewable-generation futures, choosing representative futures, identifying critical units, and solving the pruned scheduling problem. The algorithms are under development as part of the WSU-ABB collaborative project. Significant progress has been made in developing the algorithms related to generation-forecasting, and initial software implementation has been developed for a case study (on wind-generation in the Pacific Northwest during a cold-front passage in October 2013). Algorithm development for the stochastic-unit-commitment module is in a more preliminary stage. Specifically, for a small-scale constructed example (based on the IEEE 14-bus model), we are pursuing implementation and evaluation of the two-stage approach to unit commitment.

3 The Renewable Generation-Forecasting Module

This section details the first module in the operational concept, an influence-model-based tool for wide-area renewables forecasting for the day ahead. The influence-model-based approach is first motivated (Section 3.1). Then, the blocks in the renewables-forecasting module are described, with a focus on the blocks related to wind generation (Section 3.2). Throughout the development, a case study of wind generation in the Columbia River Gorge of Washington/Oregon on September 22, 2013, is used to illustrate the model.

3.1 Why an Influence-Model-Based Approach?

Many methods have been developed for renewable-generation forecasting, which span multiple temporal and spatial scales. Broadly, renewable-generation forecasting approaches can be classified based on their temporal resolution and look-ahead, spatial resolution, underlying modeling mechanism (physics-based vs. empirical), and their ability to capture uncertainties, among other factors (10). Day-ahead unit commitment requires models with look-ahead horizons of 24-48 hours, preferably with temporal resolutions of 5-15 minutes (which is fine enough to capture hourly generation profiles in some detail). Additionally, the models must be able to provide relatively accurate predictions of generation from wind farms, as well as for solar farms and/or distributed solar generators in a locality. At the same time, the models must be able to provide predictions of such generation across the geographic domain covered by a transmission system operator. Finally, the proposed approach to generation scheduling requires stochastic generation futures. While many deterministic models for generation-forecasting for wind-farms at the time-resolution of interest have been developed, fewer models capture uncertainties and are extensible to wide-area prediction. Among the models that do capture uncertainties, many simply identify error bounds around a nominal forecast, and hence do not naturally provide futures or trajectories of generation.

Among the models in the literature, ensemble-forecast-based approaches are the most relevant to generation-forecasting for day-ahead unit commitment. These approaches translate commercially-available ensemble forecasts for environmental parameters (winds, humidity, etc.) into multiple generation futures. The approaches are relevant and appealing for wide-area forecasting, in that they 1) have the proper look-ahead horizon for forecasting (typically, up to three days), 2) are able to provide predictions across a wide area, and 3) directly yield stochastic futures. The models are also appealing in that they use physics-based representations of environmental processes, and in that they are available in the public domain. However, we believe that the ensemble-forecasting approaches cannot be used directly for day ahead resource scheduling, for several reasons:

1) Ensemble forecasts typically have a temporal resolution of 3 hours at a one-day look-ahead, and a spatial resolution of 15-40 km. The temporal resolution is insufficient for day-ahead scheduling. Likely, environmental conditions may vary sufficiently across a forecast grid square (particularly in complex-terrain regions) to reduce forecast accuracy. Thus, higher-resolution forecasts are needed, particularly in geographic regions with a high density of wind and/or solar generators. From another viewpoint, interpolation of the ensemble forecasts in both space and time is needed.
2) Ensemble forecasts only capture uncertainties in initial conditions. However, wind and solar generation often may be significantly impacted by uncertainties at shorter temporal/spatial scales, which are not forecasted. For instance, beginning and end times for wind events are often highly uncertain even at short time horizons. Likewise, on partly-cloudy days, insolation on solar panels may exhibit significant short-time-scale fluctuations and uncertainties. We note that these smaller-scale fluctuations may exhibit significant temporal and spatial correlation. While a generation-forecasting tool need not capture these smaller-scale patterns exactly, it should be able to account for the resulting uncertainty to some extent.
3) Ensemble forecasts typically only provide a small number of potential weather futures, and hence generation futures. Scheduling potentially may require a larger number of possible generation futures, or at least a representative set that better spans the space of possibilities.

The influence model [14-17] is promising for addressing these needs while still leveraging ensemble forecast products, and hence can provide effective generation forecasts for unit scheduling. Specifically, the model naturally permits simulation at a desired temporal and spatial resolution (including at multiple scales across a region), while matching ensemble forecast probabilities at snapshot times and locations. Additionally, the influence model – which is a stochastic-automaton model – does capture complex uncertainties and patterns in weather evolution, in a way that permits tuning of spatial and temporal correlations. Finally, the model is simple enough to permit rapid simulation and some statistical analysis of many generation trajectories. It is worth noting the influence model has been used to model environmental uncertainties and their impacts in the transportation domain (see [14-16]); this work pursues development of analogous capabilities for generation unit scheduling.

The influence modeling approach to generation forecasting is suited for the modern control room, which has pervasive cyber- technologies. As discussed below, the prediction tool leverages current ensemble forecasting products: the modern control room is designed to access high-volume data (such as ensemble forecasts) through the Internet, and would have the capability to use up-to-date weather data as required in the proposed approach. Additionally, the approach exploits the simulation, analysis, and visualization capabilities available in the modern control room. Thus, it holds promise to provide operators with new, information-rich decision-support and automation for planning under uncertainty.

3.2 Module Blocks: Overview and Details

The renewable-generation-forecasting module involves two parallel tracks, one of which simulates possible wind-generation futures and the second of which predicts solar generation futures, for day-ahead resource planning. Here, only the blocks associated with the wind-generation track are discussed. Details on the solar-generation-forecasting track can be found in the companion paper [17]. To begin, let us note that the generation-forecasting-module, as a whole, outputs representative futures of wind and solar generation for each bus in the studied power-system model for the day ahead. To develop these futures, the module uses current ensemble forecast products, as well as archived data (on wind-farm locations and compositions, historical generation profiles, regional solar-generation usage, etc). The blocks comprising the module are envisioned as being implemented in software in the TSO's control room, for use in day-ahead planning. Here, each block's functionality is discussed, and prototype software implementations are illustrated.

3.2.1 Data Extraction Block

The data-extraction block is tasked with downloading weather-forecast data from online ensemble forecast products on a daily basis, for use in renewable-generation forecasting for the day-ahead market. A range of ensemble forecast products are posted on-line in real time, many of them by the United States National Oceanic and Atmospheric Administration (NOAA) and by European counterparts ECMWF [36]. The various ensemble forecasts each use high-resolution deterministic physics-based models for atmospheric dynamics. Multiple ensemble members or futures are produced through randomization of uncertain model parameters and/or initial conditions, with most forecast products including 15-30 ensemble members. The full models are extremely high dimensional and time-consuming to run, usually requiring several hours on a large cluster. Only a subset of the model's states variables are posted to the online server, at a moderate spatial and temporal resolution. Even this lower-resolution filtered output is quite high dimensional, typically requiring tens of gigabytes for storage. The model data on the NOAA servers can be further filtered by the user prior to downloading, permitting extraction of only relevant weather parameters in the geographic region of interest. Specifically, the data can be accessed and parsed via unix script commands. The data is encoded in the *grib2* format, which is commonly used for environmental data sets. Once downloaded, the data can be automatically translated into other common data formats (e.g., csv, plain text, etc.), again using unix scripts. Alternately, several grib2 data readers are available on NOAA's webpage [33], which can be used for display and manual processing of the data.

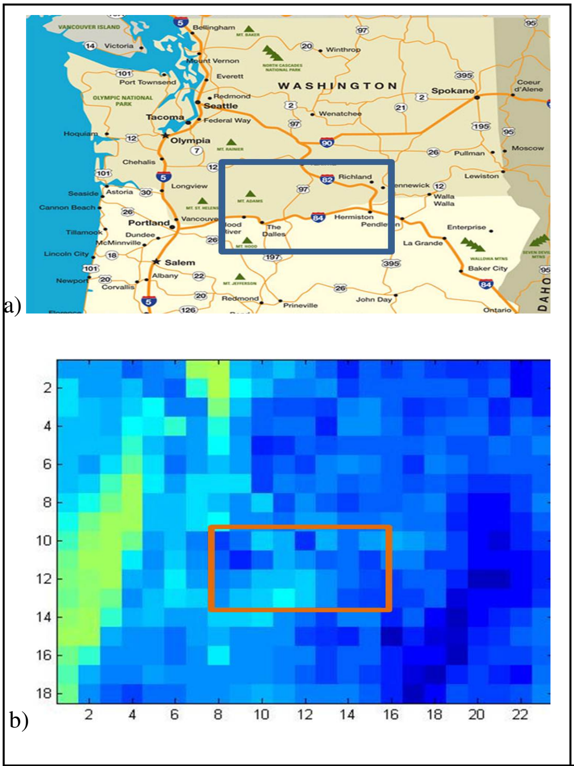

Fig. 2a) A case study of generation forecasting in the Columbia River gorge area of Washington and Oregon is considered; the boxed region has a high density of wind farms, and is the focus area of this case study. b) Snapshot wind map of one ensemble member in the Short Range Ensemble Forecast, with the region of interest indicated. Darker shades of blue indicate low wind speeds, while stronger shades of lime indicate high wind speed.

In the proposed solution, the data-extraction block is responsible for extracting a small subset of the environmental parameters needed for generation-forecasting over the time-horizon of interest, downloading this data to the TSO's local server, and translating it into a convenient form for further processing.

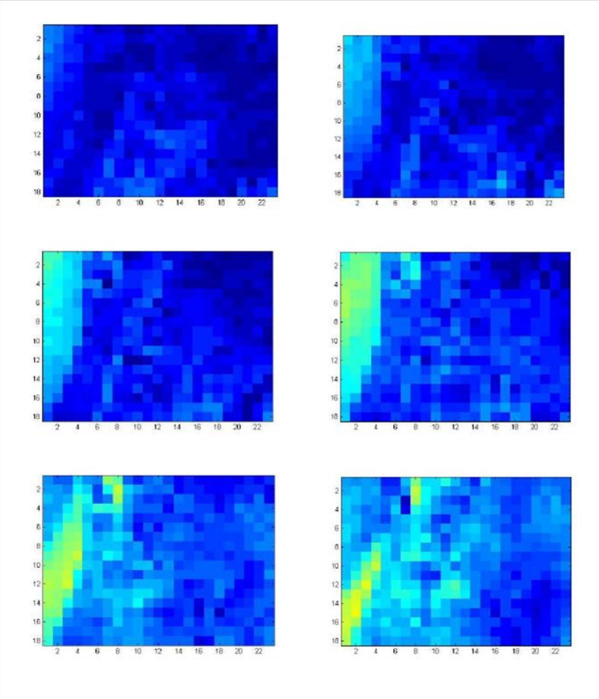

Fig. 3 The time-progression of one ensemble member in the SREF is shown. Specifically, wind speeds predicted by the ensemble member at 3hr intervals across the Pacific Northwest are shown. A cold front is encroaching on the area, leading to a period of increased winds.

Specifically, for wind-generation fore-casting, wind speed and direction vari-ables just above ground (10-30 m) are needed, across the geographic region managed by the TSO and over the full day ahead (24 hours, from one midnight until the follow-ing). Solar-generation forecasting is typi-cally more compli-cated, using temp-erature, humidity, cloud cover, and possibly other fore-cast data, see [17] for details. The data-extraction block is tasked with downloading this data to the local server, and formatting for use by downstream software blocks.

In this study, the 40km Bias-Corrected Short Range Ensemble Forecast (SREF) for the Continental United States (CONUS) is being used as the forecast data source for the generation-forecasting module. The SREF is appealing for generation scheduling in that the appropriate weather parameters (wind speed, wind direction, etc.) are posted at sufficient frequency (every 6 hours), over an appropriate look-ahead horizon (up to 87 hours into the future) and resolution (3 hrs temporal resolution, 40kmx40km grid squares for the spatial resolution), and with an acceptable delay (forecast becomes available about 2 hours after the initial forecast time). We have chosen to use the 40km bias-corrected model, specifically, because there has been an extensive effort to validate this model version. New higher-resolution versions have recently become available, and may be practical for use in the near future. To permit exploratory study, the DeGrib tool is being used to process SREF data, although basic scripts for automated downloading and processing have also been written. We note that the SREF forecasts are produced at 3Z (3AM Zulu Time), 9Z, 15Z, and 21Z daily: the TSO would use the most recent available forecast, which depends on the TSO's schedule for resolving the day-ahead market. For most markets in the CONUS, this is the 9Z, 15Z, or possibly 21Z forecast.

Case Study: Wind Generation in the Columbia River Gorge

For the case study, wind speed and direction variables were extracted from the 21Z forecast on 9/21/2013, for the period between 6Z and 21Z on 9/22/2013. Based on our focus on generation in the Columbia River Gorge region, the forecast data was extracted for the Pacific Northwest region of the United States (the states of Washington and Oregon, and adjacent areas in Idaho and in the Pacific Ocean). In Figure 2, this region is shown, and the wind speed forecast for one ensemble member at a snapshot time is also displayed; the Columbia River Gorge region is encircled in orange for convenience. Figure 3 compares the wind-speed forecasts at different times for a particular ensemble member, while Figure 4 compares two ensemble members at a snapshot time. During the period of interest, a cold-front is encroaching on the

Pacific Northwest, leading to an increase in wind speeds over the period. While different ensemble members are generally similar, they show noticeable differences in cold-front timing and strength, leading to significant variability in wind speed profiles in the Columbia-River-Gorge region.

3.2.2 The Influence Model Builder

The next blocks in the renewable-generation-forecasting module are tasked with building influence models for wind and cloud cover *(Influence Model Builder)*, which are then used to simulate wind- and cloud- cover profiles (*Influence Model Simulators*), and hence to forecast generation (*Wind-to-Generation Translator* and *Cloud-Cover-to-Generation Translator*). Here, the wind influence model is described, and the algorithm for building specific instances using the ensemble forecasts obtained from the data-extraction block is overviewed. The analogous influence model for cloud cover dynamics is described in [17].

Broadly, the **influence model** is a networked-Markov-chain or stochastic-automaton-network model, which tracks the evolution of discrete statuses across a network of interacting sites. Each site's status evolves in a Markov fashion, via simple interactions with neighboring sites. The model is appealing in that update rule is simple enough to permit rapid simulation and statistical analysis, yet can capture complex spatiotemporal evolution patterns and correlations. The model was originally envisioned as a representation for failure propagation in complex networks [15], but subsequently has been used to model e.g. inter-personal communication patterns, decision-making in sensor networks, and convective-weather evolution, e.g. [20,40].

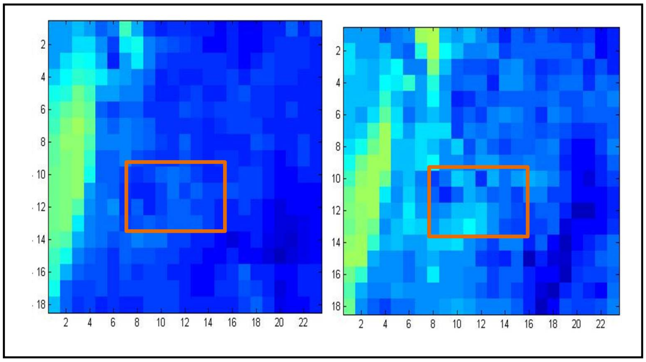

Fig. 4 Two different ensemble members are compared at a snapshot time. Although the predicted wind pattern is similar overall, there is significant variation in the region of interest.

Here, an influence model is considered that forecasts discrete wind levels in *N* contiguous sub-regions across an area of interest, at an appropriate temporal resolution for decision-making (e.g., 15 minutes) over the day ahead. Specifically, at each time step, each subregion $i \in 1,..,N$ is modeled as being in one of m statuses (labeled $1,..,m$), which identify different wind speed and direction bins (intervals). For instance, the model may use *m=6* bins to identify wind speeds between 0 and 30kph, in bins of 5kph each. Alternately, if wind speeds and directions are both tracked, *m=6x4=24* bins could be used to capture the same wind-speed levels as well as the wind-heading quadrant. The status of subregion *i* at time *k* is denoted as $s_i[k]$. These statuses evolve with the time step *k* based on a simple, Markovian update rule. This update rule captures that forecast wind characteristics in a subregion follow a statistical distribution (as extracted from the ensemble forecasts), but also show persistence over time as well as correlation across space. Specifically, the next-status $s_i[k+1]$ of subregion *i* is determined via the following two-stage update:

1. Geographical neighbors of subregion i (including the subregion i itself) are viewed as influencing the next status. To define this influence, each neighboring subregion j is modeled as providing an m-element probability vector $\boldsymbol{a}_{ij}(s_j[k], k)$, which depends on its current status $s_j[k]$ and also may vary with the time step *k*. The vectors $\boldsymbol{a}_{ij}$, as probability vectors, are element-wise non-negative and sum to *1*.
2. A weighted average of the neighbors' probability vectors, $\boldsymbol{a}_i = \sum_{j \in N(i)} d_{ij}[k]\boldsymbol{a}_{ij}(s_j[k], k)$, is computed, where the weights $d_{ij}[k]$ are assumed to be nonnegative and sum to *1*. The probability vector $\boldsymbol{a}_i$ is used to realize the next status of $s_i[k+1]$ of subregion *i* (independently of all other realizations). That is, the next status is selected stochastically to be one of the discrete possibilities $1,..,m$, with the probability that the status is q given by qth entry in $\boldsymbol{a}_i$.

The time-$(k+1)$ statuses of all sub-regions are determined simultaneously in this fashion, and the process is repeated for each time step. The influence model update is illustrated in Figure 5.

The wind influence model, as defined above, is a stochastic automaton model that produces probabilistic futures of wind trajectories in geographical subregions within an area of interest. Of course, the futures produced by the model crucially depend on the model's parameters, namely the *local influence vectors* $\boldsymbol{a}_{ij}(s_j[k],k)$ and the *network weights* $d_{ij}[k]$. Prior to using the model, these parameters must be selected so that the model produces wind futures that reflect real environmental conditions. Here, the model is parameterized based on the wind forecasts extracted from ensemble forecast products – this is what is meant by "building the influence model".

Several of the previous studies on the influence model have considered parameterization from data or forecasts: broadly, the sparseness of the model often permits parameterization from a fairly limited data set. Most relevant to the research presented here, a method was previously developed for para-meterizing influence models for convective-weather-propagation from ensemble weather forecasts, in the case where the subregions are grid squares. This approach has been adapted to parameterize the wind influence model.

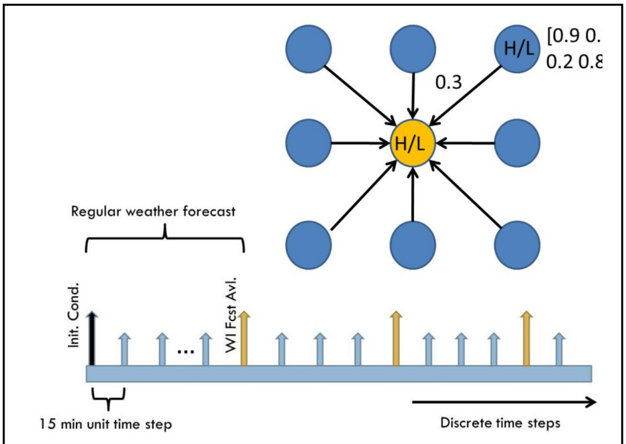

Fig. 5 The influence model update rule is illustrated (top diagram). The influence model is parameterized to statistically match the ensemble forecast at snapshot times, as diagrammed below.

Specifically, let us consider building a gridded influence model using ensemble forecast products at a certain spatial and temporal resolution (e.g., 40km grid squares and 3 hour temporal resolution for the SREF). While the ensemble forecast data has been extracted over a wide area, wind generation is often concentrated in regional clusters – e.g., in the Columbia River Gorge area in the Pacific Northwest. Here, one of these clusters is used to define an area of interest for influence modeling. Within the area, a model with higher resolution than the ensemble forecast may be needed, to capture dynamics and variations at the spatial scale of wind farms, and at an appropriate temporal resolution for unit scheduling. Specifically, a y-fold increase in the spatial resolution (along each dimension) and a z-fold increase in the temporal resolution is assumed. To build the influence model, the ensemble forecast data is first used to determine desired status probabilities for influence-model grid squares at snapshot times (for instance, every three hours if SREF data is used). Specifically, each influence model grid square is located within an ensemble-forecast grid square; the fractions of ensemble members in each status (wind speed/direction bin) in this square can be viewed as desired status probabilities for the corresponding influence model square. In this way, the ensemble forecasts provide local status probabilities across the region of interest at every z time steps. Once these snapshot probabilities have been computed, the technique for influence-model parameterization given in [14-17] can be applied directly. Specifically, the technique allows selection of the influence-model parameters so that the local status probabilities at the snapshot times exactly match the desired snapshot probabilities, and further the status probabilities at intermediate times are interpolations of these desired probabilities (see Figure 5). In addition, the parameterization technique gives the user the freedom to tune the extent of spatial and temporal correlation (or persistence) in the wind profile. In these initial studies, these correlation-tuning parameters have been chosen so that wind deviations have significant persistence for about one hour and 100km. In the future, we expect to tune the parameters to match historical correlations in wind speeds/directions at wind farms. Since the technique for parameterizing the influence model was presented in earlier work, it is not described in detail here, see [17] for these details.

A couple of remarks about the influence modeling approach are worthwhile. First, the reader will note that the model uses a binned or discretized representation of wind. An alternate continuous-valued influence model can be envisioned [37]. However, we believe the binned approach is appropriate because of intrinsic precision limits (and limited precision needs) for the day-ahead wide-area forecasting goals of this project, and because binned models naturally permit translation to generation levels (see Section 3.2.4). The bin resolutions can be chosen at the user's discretion, so varying levels of precision are possible. Second, while our focus here has been on building a gridded model with a single resolution, the influence model permits arbitrary subregion topologies. The parameterization (model-building) technique also can be extended to more general topologies. A multi-resolution gridded model is currently under development.

Case Study

The influence model builder has been implemented for the described case study. Specifically, an influence model for wind speeds has been developed for the highlighted area of interest, corresponding to the Columbia River gorge region which has a high density of wind generation. Noting that the wind speeds during the period of interest are low (less than 12 miles/hour), a coarse binning model has been used in this initial study. Specifically, two wind-speed-based bins are assumed, one corresponding to wind speeds below 6 miles/hour (which permits no wind generation for most turbines) and the other to wind speeds between 6 mi/hr and 12 mi/hr (which is above the cut-in speed but only yields a low level of wind generation). The SREF forecast was used to build the influence model simulator, as discussed above. The implemented influence model achieved 12-fold multiplication in the temporal resolution (to provide forecasts every 15 minutes) and a 2-fold multiplication in the spatial resolution (yielding 20kmx20km grid squares) compared to the ensemble forecast.

3.2.3 Influence Model Simulator

The built influence models for wind and cloudiness characteristics can be used to simulate or produce a large number of wind/cloudiness futures over the day ahead. The software tool that does this is referred to as the influence model simulator. The influence model simulator is discussed here, with a particular focus on wind-future simulation (see [17] for a discussion of cloud-cover simulation for solar-generation prediction).

Since the wind influence model is a stochastic automaton network model, each simulation of the model yields a different future or profile. Specifically, the influence model's update rule (see Section 3.2.2) is applied over the modeled time horizon (the day ahead), for the built model. Simulating the model in this way produces a specific wind bin profile at the specified temporal and spatial resolution. By repeating the simulation many times, a large number of independent profiles or futures is obtained. These futures are each different, but their aggregate statistics match the designed statistics of the built influence model (including local status probabilities and correlations), and hence also match the ensemble-forecast statistics at snapshot times.

The influence model update rule permits fast simulation, only requiring computation of a linear function followed by a randomization (which can be achieved by producing a uniform random variable on [0,1] and comparing it with a threshold). The simulation time scales linearly with the number of grid squares and the number of time steps simulated. For a realistic-scale model (say, 100-5000 grid squares over a full day), thousands of futures can be produced in less than a second. The special structure of the wind influence model also permits efficient statistical analysis, including characterizations of temporal and spatial correlations, as well as variability in aggregate wind characteristics.

Case Study: The influence model simulator has been implemented for the Columbia River gorge case study. Specifically, the built wind influence model has been used to produce 1000 wind futures. Snapshots of two futures are shown in Figure 6. Both futures show a trend toward increasing wind speeds, reflecting the trend in the ensemble forecast. They also show certain common spatial characteristics (for example a consistently low wind speed in a couple of the Southern grid squares), which reflect topological impacts on wind characteristics. However, the two futures show considerable variability in the wind speed profile, and also display complex spatial patterns. In Figure 6, we also map the probability of the higher-wind-speed bin at the snapshot times. The influence model simulations match these local status probabilities in aggregate, but show considerable variability and also enforce spatial and temporal correlation.

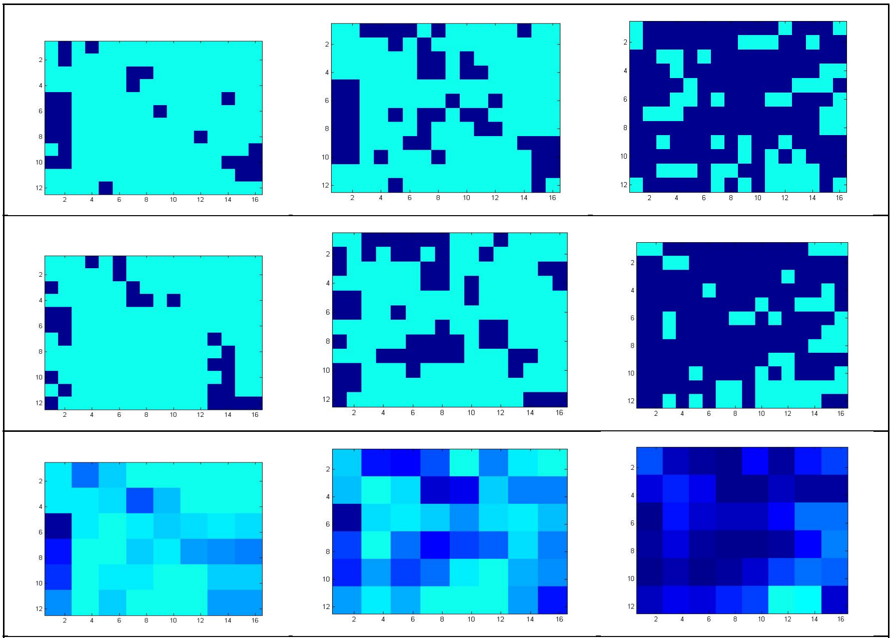

Fig. 6 Top *and middle rows:* two simulations of the influence model are shown (darker squares represent the higher-wind bin, lighter squares indicate the low-wind bin), for the Columbia River Gorge region (the boxed area on the ensemble forecast). This influence model has twice the resolution of the ensemble forecast. We notice that the influence model captures the spatiotemporal trends in wind speed in the ensemble forecast, but captures significant smaller-scale variations and correlations in wind, particularly during periods of rapidly-changing weather. *Bottom row:* the probability of the higher wind-speed bin across the region of interest is shown at snapshot times, as obtained from the ensemble forecast.

Finally, in Figure 7, we have presented some aggregate statistics of the wind-influence-model simulations (e.g., a histogram of the total number of high-wind grid squares), to illustrate the level of variability among the futures produced by the model. These statistical analyses show that the variability changes significantly with time and location, with highest variability when expected wind speeds are changing rapidly. The standard deviations in wind generation found in this way roughly match the uncertainty levels given in the literature [29, 34].

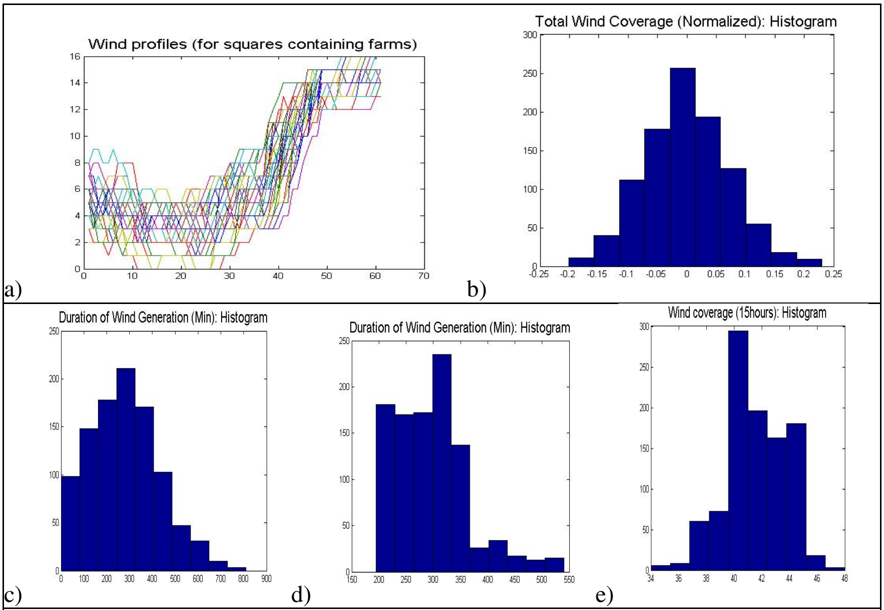

Fig. 7 Statistical characterizations of wind and wind generation profiles are obtained, using the influence-model-based simulator. a) The number of grid squares with high wind is traced, for multiple influence-model trajectories. b) A histogram of the total number of high-wind periods at all wind-farm locations is shown (normalized to the mean); the variability is about 10%, which matches with observed data. c) and d) Histograms of the time duration of wind generation (i.e., the total time such that winds are high enough to generate power) are shown, for two different wind farms. The distributions are quite different. e) The distribution of total wind coverage across the region is shown, at a snapshot time; there is high variability at times when the average wind speed is changing rapidly.

3.2.4 Wind- to Generation- Translator

The next blocks in the generation-forecasting module are responsible for translating the wind futures produced by the influence model to wind-generation futures (and analogously to translate cloud-cover futures and other environmental parameters to solar-generation levels, see [14]). There is a wide literature on modeling wind turbines and wind farms, which can be brought to bear to forecast wind generation from wind profiles [10-12]. Unit commitment requires forecasting at the resolution of wind farms across a wide area, rather than precise forecasting at a single location. Two techniques are worth reviewing. First, *binning techniques* have been used to model wind generation at the level of wind farms. In these techniques, wind bins (which may involve both speed and direction parameters) are mapped to wind generation levels for a specific wind farm, using statistical analyses of historical data. The binning approaches dovetail nicely with our solution, since the influence-model-simulator produces binned wind futures in subregions across the wide area. For a wind farm located in a particular subregion, the forecast wind bin level in this subregion can then be translated to a generation level using the binning-based model. Using maps of the wind farms' connections to the electric power grid, the wind generation at each bus during each time step can be determined. By applying this method to each influence-model-produced future, many wind-generation futures can be obtained.

Alternately, simple physics-based models for a wind turbine can be used for generation forecasting. The simplest models approximate a turbine's wind generation as a cubic function of the wind speed, between a lower cutoff speed and the turbine's rated wind speed; the model's parameters depend on the type of wind turbine being modeled. To translate the influence model futures via this model, we again determine the wind bin level at the wind-farm location of interest from the influence model. This bin level for the wind farm then is converted to a single wind speed : either the median wind speed in the bin may be used, or the speed may be randomized within the bin (with the motivation that very small scale variations in speed cannot be captured and may be modeled as uncertain). The physics-based model can be used to determine the wind generation for each turbine in the farm, and hence total farm-level generation can be determined. The remaining procedure for obtaining wind-generation futures is the same as for the first approach. We note that much more intricate models for wind turbines and farms are available, that account for wind-direction effects, capture wake effects and topographical variations, etc. However, noting that wide-area forecasts are needed, these simple approaches are approaches.

Both approaches for wind- to generation- translation described above require some data on wind farms and turbines. For both approaches, the locations of wind farms to be modeled must be known. For the binning approaches, historical data on wind speeds and generation for each farm is also needed, so that binning-based models can be constructed for each farm. These models also must be updated if generation capacity is added to a farm, and when new farms are brought online. For the physics-model-based approach, the number of wind generators of various types must be known for each wind farm. ISOs and TSOs typically have available to this information (see e.g. a discussion of wind farm data for the Texas grid [35]).

3.2.5 Representative-Future Selector

The final block in the generation forecasting module is tasked with selecting a few representative wind- and solar- generation futures, from the large set of simulated futures. Representative futures are sought both for use in unit scheduling and to provide operators with concrete illustrations of wind profiles for decision making. More specifically, our motivation for selecting a subset of futures as representative ones is three-fold: 1) to provide operators with an indication of the range of weather outcomes and consequent wind/solar generation profile that may occur on a given day, 2) reduce the computation needed for the unit-commitment problem (albeit in a somewhat limited way, see Section 4 for a discussion), and 3) facilitate performance evaluation of the unit-scheduling design by identifying typical test cases for weather outcomes.

Broadly, the purpose of the representative-future-selection block is to 1) choose a sparse set of futures that span the range of wind/solar-generation outcomes for the day of interest and 2) assign likelihoods to these typical futures. Several approaches have been developed for selecting representative samples of a random variable or random process. In this project, a technique for sample selection known as the *Probabilistic Collocation Method (PCM)* [20] is considered, which draws on a numerical-integration method known as Gaussian quadrature. This method allows selection of futures according to one or more selected performance measures (specifically, to span the range of possibilities for these performance measures). Thus, the algorithm automatically selects futures that are ordered with respect to the performance measures. If appropriate performance measures are used, the method can be used to distinguish futures which will require significantly different commitment and dispatch plans, as is needed for stochastic unit commitment. Since PCM has been developed in previous work, technical details are omitted. For this project, a Matlab software implementation of the method has been developed as part of the generation-forecasting module.

One key challenge in using the proposed approach is to choose appropriate performance measures to select representative futures. To be useful, performance measures must be able to distinguish weather/generation profiles which will require significantly different schedule profiles. In this first effort, we use as the metric the total wind generation, which should be strongly indicative of conventional generation requirements and hence these units' schedule. More broadly, we anticipate relying on operators' experience to choose performance metrics. In addition, historical data can be used to regress possible performance metrics against schedule profiles, to determine how predictive the measures are of the schedules.

Case Study

The representative-future selector has been implemented for the Columbia River gorge case study, using the total wind generation as the metric for selection. Specifically, the analysis considers generation from six large wind farms in the Columbia River Gorge area, which are connected to two buses in an example power-system model (see Section 4.3). This wind-farm example is constructed, but the locations and sizes of the farms are similar to those of actual farms in the Gorge area. For each wind future produced by the influence model, the wind-power generation at the two buses of interest is computed, using the simple physics-based model for wind-to-power translation for a wind turbine. PCM has been applied to obtain five representative futures (very low power, low power, medium power, high power and very high power generations) from 1000 wind power generation futures, using the total wind generation at the two buses as a metric for future selection. Figure 8 and Figure 9 show four of the representative futures (with Figure 8 showing the generation at Bus 1 and Figure 9 at Bus 2). The very low power generation future is not shown here, because the generation level is nearly zero. Both figures show increasing wind power generation trend, which is consistent with the weather forecasting.

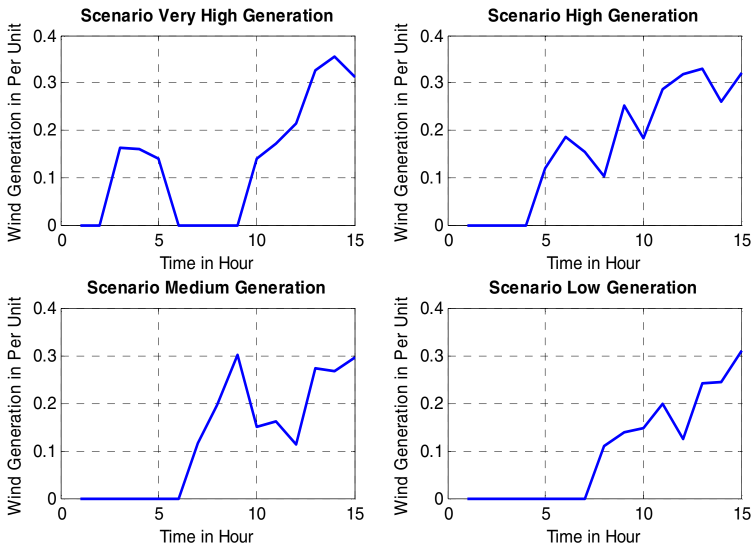

Fig. 8 Representative generation futures, Bus 6

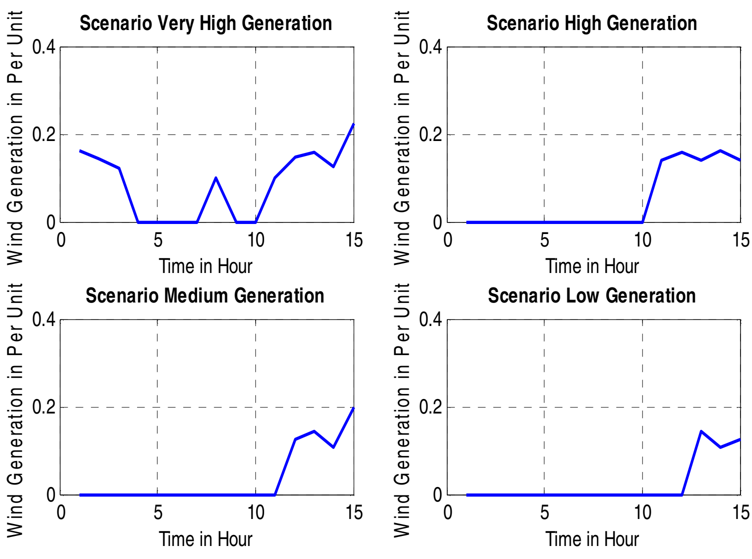

Fig. 9 Representative Futures, Bus 8

The PCM tool also assigns a probability for each scenario. These probabilities are shown in Table 1, to illustrate the likelihood of each possible representative generation future. These statistics show that wind power has more than 95% probability to be in low, medium and high generation levels. Meanwhile, the extreme cases (zero generation and very high generation) have less than 3% probability. The probability distribution of the representative futures statistically match the large ensemble of futures produced by the influence model.

Table 1 Likelihoods of the representative generation futures (scenarios)

Scenario	**Probability P_i**
Very High Wind Power Generation	0.0139
High Wind Power Generation	0.1314
Medium Wind Power Generation	0.4658
Low Wind Power Generation	0.3600
No Wind Power Generation	0.0289

4 Toward a Scheduling Module: Some Initial Explorations

The proposed end-to-end solution for day-ahead resource planning requires implementation of a unit commitment or scheduling algorithm, which is the focus of the scheduling module in Figure 1. This scheduling module is responsible for selecting an on/off schedule and dispatch plan for conventional generator units for the day-ahead market, which accounts for uncertain generation from the (non-dispatchable) intermittent-renewable units. Specifically, the scheduling module is tasked with using the representative wind- and solar- generation futures outputted by the generation-forecasting module, along with knowledge of the power grid and the market, to design unit schedules and dispatch levels for the conventional generators. At its essence, implementing the module requires solving a *stochastic unit commitment* problem. Our focus in building the scheduling module is to use a stochastic unit commitment algorithm that is practical for implementation in the current transmission-system operational paradigm, rather than to propose a new decision-making paradigm. Our perspective is to view stochastic unit commitment as a two-step process, first requiring an identification of critical conventional units whose schedules are dependent on the renewable units' generation futures, and second achieving an optimization of these unit's schedules and hourly economic dispatches.

The development of the scheduling module is still a work in progress, and here only some preliminary explorations are presented. Specifically, the wide literature on stochastic unit commitment algorithms is briefly reviewed, and challenges in integrating these algorithms into control-room software technologies are discussed (Section 4.1). A simple example problem is then introduced (Section 4.2), and then used to illustrate the scheduling module (the selection of critical units and the optimal scheduling of these units) in Section 4.2. The unit-commitment module is described in detail in the context of the example (Section 4.3). Finally, design results for the example problem are summarized (Section 4.4), and a performance evaluation of the model is undertaken (Section 4.5). It is important to stress that these preliminary explorations ignore many features of importance in stochastic unit commitment (e.g., security constraints), and certainly should not be interpreted as achieving a complete solution. Instead, these explorations expose subtleties in developing unit commitment plans across renewable-generation profiles, and illustrate the proposed two-step process.

4.1 Related Literature on Stochastic Unit Commitment

Unit commitment (UC) refers to the on/off scheduling as well as hourly dispatch of available generation units over a planning horizon (often, the full day ahead) to meet the time-varying electric load. Most TSOs routinely use UC software for resource planning, most commonly for dispatch in the day-ahead market. Typically, the UC plan is obtained by solving a deterministic optimization problem, to achieve a lowest-cost scheduling and dispatch of conventional generation. The growing penetration of intermittent renewables, which have significant uncertainty at a one-day look-ahead, is creating challenges to system operators to manage load/generation balance. Thus, there is a strong motivation to develop new unit commitment algorithms that allow scheduling/dispatch of conventional generation while accounting for uncertainty in renewable generation.

A number of methods for unit commitment under uncertainty have been developed in the literature (many of them focusing particularly on wind-generation integration), under the headings of stochastic unit commitment (SUC) and stochastic security constrained unit commitment (SCUC). These papers broadly focus on the problem of scheduling and dispatching generation to optimize an expected cost in the face of generation/load uncertainty, but vary significantly in 1) modeling generation/load uncertainties, 2) the cost function and specific design problem, and 3) the methods used for optimization, among other differences (30-32). Several recent works by Oren's group and others as being particularly aligned with the approach pursued here, in that they consider scheduling given multiple stochastic futures of wind generation. The study of Constinecu et al on exploiting ensemble forecasts for stochastic unit commitment [38], and the efforts of Sauer and his co-workers on uncertainty management (e.g. [39]), are also closely aligned with the research described here.

Although these stochastic SCUC models are promising tools to solve the unit commitment problem with large scale renewable energy integration, to the best of our knowledge they have not yet being used by TSOs in the control room. The perspective of this chapter is that advances in several directions are needed integration of these approaches into control-room technology. First, realistic models for wind/solar generation that leverage weather-forecast products are needed within the stochastic unit commitment solutions. Many existing studies make simpler assumptions regarding wind profiles, for example [21,28] and derivative works use Monte Carlo simulation to generate wind speed and assume the wind speed error distribution is Gaussian. Second, the unit-commitment strategies need to be tailored to permit easy implementation in the current operational paradigm. In particular, many of the stochastic unit commitment approaches assume hourly re-planning of the commitment plan per a dynamic-programming solution, but most TSOs use a binding day-ahead market and hence require a fixed optimal plan. Additionally, a practical implementation would benefit from a performance evaluation of the designed commitment plan over the possible weather futures, and simple display of plan specifics and performance characteristics. Third, stochastic unit commitment remains computationally challenging for problems of realistic scale (1000's of buses, 100's of generators), and further techniques for reducing problem complexity are needed.

4.2 An Exploratory Example

The scheduling module has been developed in the context of a small-scale example, based on the IEEE 14-bus test system. The example system is assumed to have both intermittent renewable generation and conventional generation. Specifically, buses 1, 2 and 3 are connected to conventional generators, while bus 6 and bus 8 connect wind generation in the Columbia River gorge area. Figure 10 shows the IEEE 14-bus test system. In this example, the wind generators can provide up to 40% of the total power. For the day of interest, the wind speeds are relatively low and, thus, the expected total wind generation is lower than 20%. It is assumed that the wind generation cannot be scheduled or dispatched, i.e. their generation levels are determined entirely by the wind profile. Our goal is to find the optimal commitment strategy (on/off schedule and dispatch) for the conventional generators. We approach the unit-commitment problem in two steps, first focusing on selecting critical units whose on/off profiles may significantly depend on the uncertain wind generation, and second solving the unit commitment problem using a pruned decision space.

Both the critical-unit selection and the unit commitment optimization require solution of hourly economic dispatches (ED). While use of ED is commonplace in power-system operations, it is useful to briefly introduce the specific ED problem considered here. The goal of the ED considered here is to minimize the total generation cost and real power losses, subject to transmission and operational constraints. For simplicity, a DC power flow is considered, with. the real power loss cost approximated as a penalty cost. Specifically, the objective function is

$$F = \min\left(C_{cg} + C_{wg} + \lambda P_{loss}\right)$$

Subject to:
Generator constraints

$$P_{min} \leq P_g \leq P_{max}$$

DC power flow line limits

$$-P_{ij} \leq \frac{\delta_i - \delta_j}{x_{ij}} \leq P_{ij}$$

And Real power balance

$$\sum_{i=1}^{n_g} P_{Gg,i} + \sum_{k=1}^{n_w} P_{Gw,k} = P_{load}$$

where n_g is the number of conventional generator and n_w is the number of wind (renewable) generators. The goal of the ED is to design dispatch level $P_{Gg,i}$ $(i = 1,2, \ldots, n_g)$ for the conventional generators. The conventional generator's operational cost C_{cg} is assumed to be quadratic in the power generation; the wind generator's operational cost C_{cw} is assumed to be proportional to the wind power generation P_{Gw}, which is obtained from the influence model and wind-to-generation translator. The real power loss is approximated as a linear function of the generation vector, in the standard way:

$$P_{loss} = (B^{-1}P)^T G^* (B^{-1}P) = P^T (B^{-T} G^* B^{-1}) P$$

where P is a column vector of real power generation, B is the network susceptance matrix and G is the network conductance matrix.

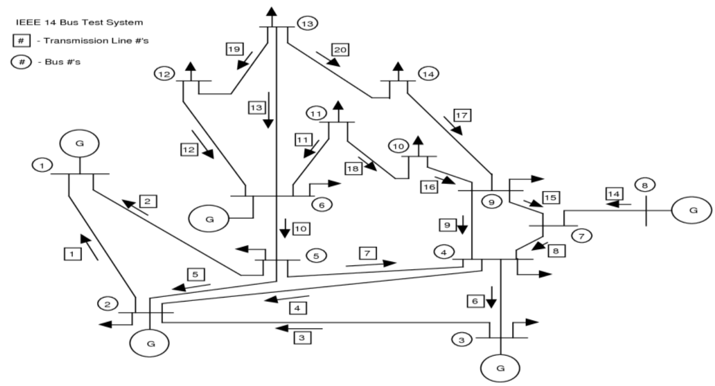

Fig. 10 IEEE 14-bus system

4.3 Scheduling Module: Details and Simulation

Broadly, the scheduling module has two stages. The first stage identifies critical units, and the second stage finds the optimal on/off schedule of critical units and the dispatch levels of all units (thus taking advantage of the critical-unit identifier to prune the decision space). Let us begin with brief descriptions of each stage's functionality, in the context of the example. It is worth remarking that the proposed algorithms leverage only the representative generation futures, which are expected to provide sufficient coverage of the uncertainty space to permit decision. In fact, all futures could be used for scheduling with a relatively modest increase in computational cost (roughly linear in the number of futures). This alternative can be implemented in an entirely analogous way.

4.3.1 Critical Unit Identifier

The critical unit identifier determines a small set of conventional (dispatchable) generation units whose on/off schedule and dispatch may be highly sensitive to the generation profile on the day of interest. Specifically, critical units are identified as follows. For each representative generation future (five in our case), the economic dispatch is determined for each hour on the day ahead, *assuming that all of the dispatchable units are on-line*. We notice that some units may or may not need to dispatch power, depending on the representative generation future. Units whose dispatch may or may not be zero depending on the weather future are considered critical units, since they will be difficult to schedule given the uncertainty in the generation profile. In the example, when the high wind power future is considered, we notice that some units are scheduled to produce zero power at some times based on the ED. Meanwhile, at the same times in the low wind generation case, those units are scheduled to produce some amount of power. These units are identified as critical units. Specifically, in the 14-bus example, conventional generator 2 at bus 2 is identified as a critical unit by solving the optimal dispatch problem in Section 4.2. Figure 11-1 and 11-2 illustrate the dispatch of generator 2 for the high and low wind generation futures, respectively. At each time hour, the dispatch level is determined by the ED problem solved in section 4.2. At times t=7, t=9 and t=11, generator 2 dispatches some amount of power in low wind generation case, but dispatches zero power in high wind generation case.

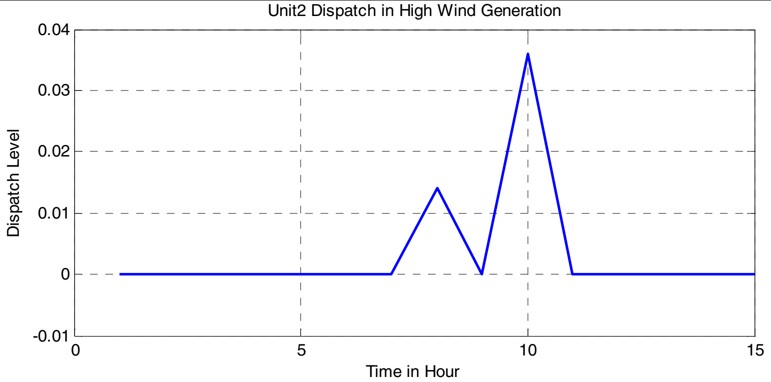

Fig. 11 -1 Unit 2 Dispatch in High Wind Generation Case

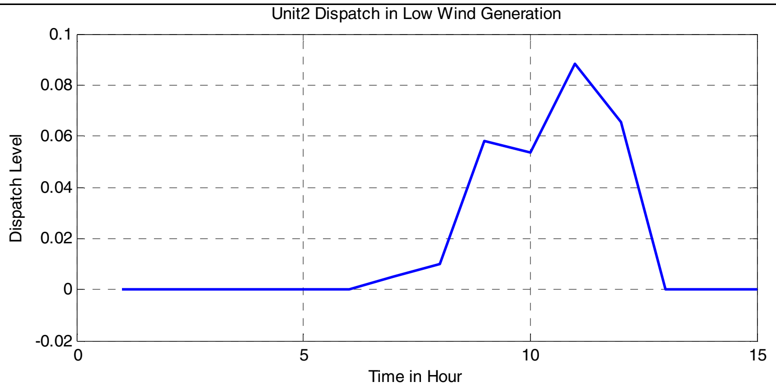

Fig. 11 -2 Unit 2 Dispatch in Low Wind Generation Case

Figure 12 and 13 demonstrate the non-critical units' dispatch levels (assuming again that all generators are on-line). Specifically, figure 12-1 and 12-2 illustrate the dispatch of generator 1 for the high and low wind generation futures respectively; and figure 13-1 and 13-2 illustrate the dispatch of generator 3 for the high and low wind generation futures respectively. No matter what the wind generation future is, the non-critical unit is assigned some nonzero dispatch level at each time hour.

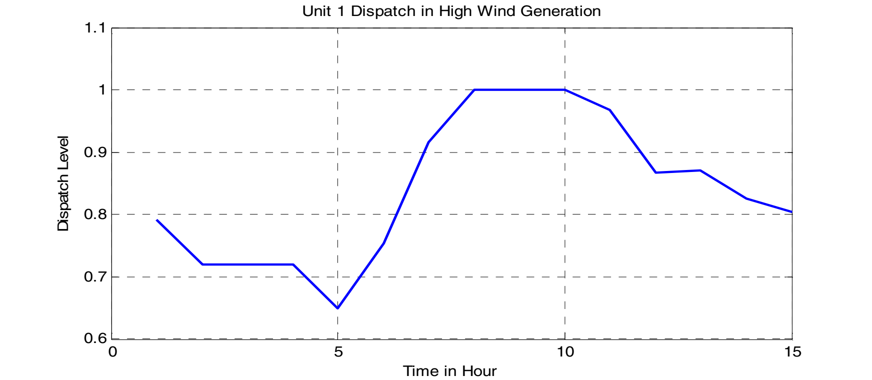

Fig. 12 -1 Unit 1 Dispatch in High Wind Generation Case

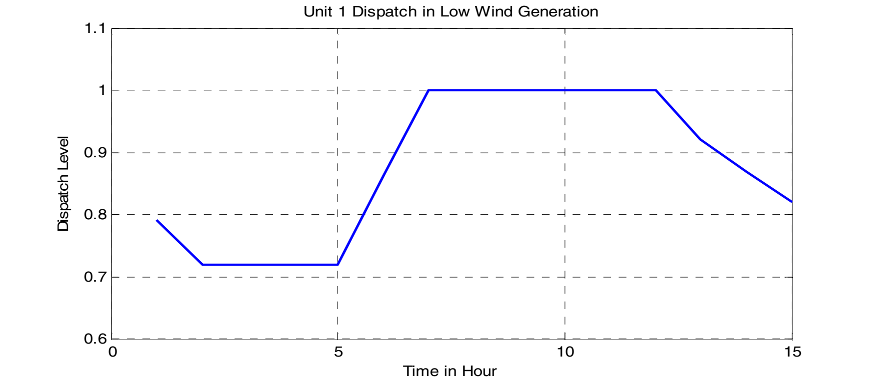

Fig. 12 -2 Unit 1 Dispatch in Low Wind Generation Case

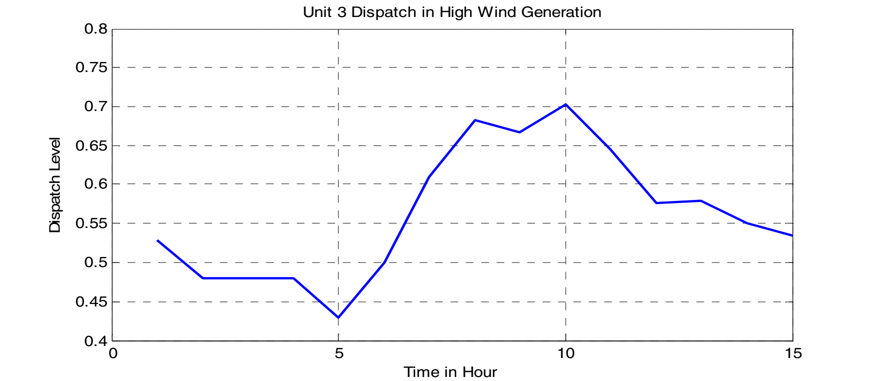

Fig. 13 -1 Unit 3 Dispatch in High Wind Generation Case

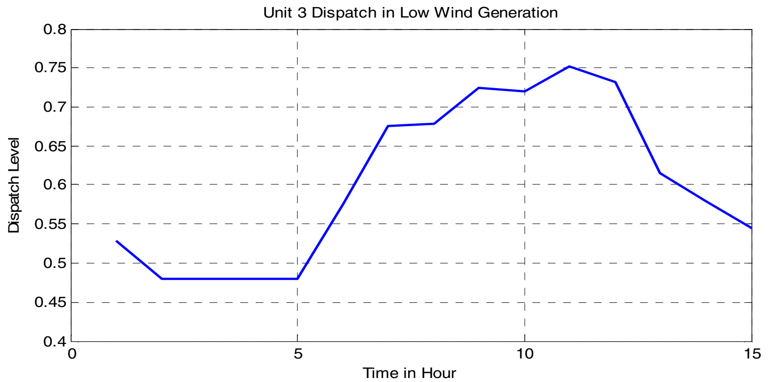

Fig. 13 -2 Unit 3 Dispatch in Low Wind Generation Case

4.3.2 Scheduling On/Off Profiles and Dispatches

The day-ahead market requires a single binding unit-commitment and dispatch schedule that performs well across weather futures. This design problem can naturally be phrased as an optimization problem, to minimize a cost or expected cost with regard to on-line/off-line schedules and dispatch level. Broadly, the optimization formulations considered can be phrased as mixed integer programming problems, which are similar in flavor to several of the stochastic unit commitment problems in the literature (although our formulation does not allow re-planning). For problems of moderate/large scale, standard optimization programs can be used to solve the unit commitment task. However, a special pruning of the design space is pursued here to reduce the high computational burden of this optimization, and to obtain insightful characterizations of uncertainty impact. Specifically, the dispatch level of all thermal units and the on/off of schedule of only the critical units are considered as design variables our formulations.

Two approaches for solving the stochastic UC have been considered. The first benchmark method uses a single mean wind power generation trajectory to design the critical units' on/off schedules and all thermal units' dispatch level. Scheduling based on mean wind profiles has been considered in the literature (see e.g. [39]), and in this sense the approach is a benchmark. The second method explicitly accounts for the wind generation uncertainty in the UC objective function, and hence seeks for a schedule and dispatch profile that minimizes the expected cost across possible weather scenarios. Let us give a mathematical description of each problem formulation.

Specifically, suppose there are n thermal units in a power system, and k of them are critical units. In method 1, the schedules of the critical units and the hourly dispatch of all units are designed to minimize the following cost:

$$f(\vec{g},\vec{x},\vec{y}) = \min(\sum_t \sum_i g_{it}C_{it} + \sum_t \sum_i y_{jt}S_{jt} + \sum_t \sum_j x_{jt}H_{jt} + \lambda \sum_t P_{loss,t})$$

Subject to the following constraints:

Power balance –

$$\sum_i g_{i,t} + E\left(g_{wind,t}\right) = P_{load,t}$$

Generation limits –

$$g_{i,min} \leq g_i \leq g_{i,max}$$

Line flow limits –

$$-g_{ij} \leq \frac{\delta_i - \delta_j}{x_{ij}} \leq g_{ij}$$

On-line time limit –

$$t_{i,online} \geq t_{i,online,min}$$

Transition costs (start-up and shut-down costs) are also modeled. They are described below after the problem formulation for method 2.

In method 2, the schedules of the critical units and the hourly dispatch of all units are designed to minimize the following cost:

$$f(\vec{g},\vec{x},\vec{y}) = \min E(\sum_t \sum_i g_{it}C_{it} + \sum_t \sum_i y_{jt}S_{jt} + \sum_t \sum_j x_{jt}H_{jt} + \gamma \sum_t (P_{L,t} - \sum_i g_{it} - g_{wt,m})^2 + \lambda \sum_t P_{loss,t})$$

Subject to the following constraints:

Generation limits –

$$g_{i,min} \leq g_i \leq g_{i,max}$$

Line flow limits –

$$-g_{ij} \leq \frac{\delta_i - \delta_j}{x_{ij}} \leq g_{ij}$$

On line time limit –

$$t_{i,online} \geq t_{i,online,min}$$

In both formulations, the decision variables are the following:

g_{it} is the power produced by the thermal generator i in time period t

y_{it} is 1 if critical unit starts at the beginning of period t, otherwise, 0

x_{it} is 1 if critical unit shuts at the beginning of period t, otherwise, 0

Index i represents all thermal units, and index j represents all critical units

Other parameters include:

$gw_{t,m}$, which is the total power produced by wind generator in time period t at representative scenario m.

Finally, the term $\gamma \sum_t (P_{L,t} - \sum_i g_{it} - g_{wt,m})^2$ in the second problem formulation is called the correction cost. This term requires some further discussion. Notice that, in method 2, the power balance cannot be enforced since the renewable-generation is uncertain. Instead, we model the imbalance between power generation and consumption as a correction cost, which reflects the additional cost needed to meet the power imbalance on the day-of-operations (through re-dispatch on the hourly market, use of reserves, and possibly through turning on fast-ramp units). We stress that the correction cost arises due to the uncertainty in real wind power generation. In contrast with other stochastic unit commitment efforts, we do not seek to model the re-dispatch of power at an hourly scale in detail, but propose the quadratic correction cost to encompass the family of corrective actions that may be taken. We anticipate that operators would choose the scaling constant γ based on historical costs incurred on the day of operations when there is a significant generation-load imbalance. In particular, a regression may be used to determine the dependence of the additional cost on the imbalance. We anticipate that the regression may identify a non-quadratic mapping for this correction cost: in particular, insufficient generation is likely to be more expensive than overproduction; we leave a careful analysis to future work.

Another point that requires discussion is the loss penalty cost in method 2. As discussed previously, we model the loss penalty cost as a function of all the thermal units' dispatch levels under the power balance constraints. In method 2, we do not have this power-balance constraint, which complicates computation of the loss penalty. Here, we assume the sum of all thermal units' dispatch level is close to the load consumption minus the mean renewable generation. Under this assumption, we can use the same penalty cost function in method 2 as is used in method 1. Finally, for method 2, we note that the expectation is computed across the representative renewable-generation futures.

The problem formulations described above have further been extended to include transition costs, including startup costs and shutdown costs. Startup costs involve both fixed costs C_f and variable costs C_v. Shutdown costs generally involve only fixed costs and sometimes are not significant. Generally, variable costs of start-up depend on two different shutdown states the unit is in. The two possible states are hot reserve and cold reserve. The mathematical model of the variable cost in hot reserve state is:

$$C_{vb} = C_b t f,$$

and in cold reserve state is

$$C_{vc} = C_c[1 - e^{-\frac{t}{a}}]f,$$

where t indicates the time period that the generator is in this state, and the remaining scalars are model parameters. Details are omitted.

How to Solve the Optimizations: The described optimization problems are mixed integer programming problems, and a range of tools for solving these problems can be brought to bear. We stress that the proposed methods take several steps to reduce the inherent computational complexity of the problems, including pre-selection of critical units (which reduces the number of integer variables) and abstract modeling of the correction cost. These simplifications are particularly important for optimization of an expected cost across arbitrary wind futures, since traditional simplifications of mixed integer programs often fail in this case.

For the small-scale case study described below, we have simply used an exhaustive search over the binary (on/off) variables together with the quadratic-programming tool in Matlab to find the optimum. For this small-scale example, it is worth noting that the optima obtained over the pruned design space are identical to those that would be obtained if all units' schedules were designed, and hence our critical-unit-based approach does not lead to performance degradation.

4.4 Example Problem: Results

Tables 3 and 4 show the dispatch levels of all units and the critical unit's on/off schedule, respectively, when the first unit-commitment method is used (i.e., the method based on expected generation).

Dispatch level:

Time	1	2	3	4	5	6	7
Gen1	0.7902	0.7188	0.7177	0.7187	0.7094	0.8486	0.9618
Gen2	0	0	0	0	0	0	0
Gen3	0.5227	0.4746	0.4738	0.4745	0.4682	0.5621	0.6385

Table 3

Time	8	9	10	11	12	13	14	15
Gen1	0.9906	1.0	1.0	1.0	1.0	0.9055	0.8568	0.8022
Gen2	0	0	0	0.0234	0	0	0	0
Gen3	0.6579	0.6915	0.7657	0.8	0.6858	0.6005	0.5676	0.5308

Table 4

Critical unit on (1) /off (0)

Time	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Gen2	0	0	0	0	0	0	0	0	0	1	1	1	0	0	0

Meanwhile, Tables 5 and 6 show the dispatch level of all units and the critical units' on/off schedules, respectively, for the second method (i.e., minimization of the expected cost over wind futures).

Dispatch level

Time	1	2	3	4	5	6	7
Gen1	0.7751	0.7053	0.7041	0.7052	0.6961	0.8322	0.9430
Gen2	0	0	0	0	0	0	0
Gen3	0.5426	0.4927	0.4919	0.4926	0.4861	0.5834	0.6624

Table 5

Time	8	9	10	11	12	13	14	15
Gen1	0.9697	0.9745	1.0	0.9931	0.9731	0.8880	0.8408	0.7870
Gen2	0	0.0051	0.0035	0.0544	0	0	0	0
Gen3	0.6841	0.7173	0.7680	0.7819	0.7181	0.6230	0.5890	0.5508

Table 6

Critical unit on/off

Time	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Gen2	0	0	0	0	0	0	0	0	1	1	1	0	0	0	0

4.5 Example Problem: Evaluation

Because the real day-ahead wind speed and power generation has large variability, operators may wish to check how an optimal UC plan performs under real weather scenario. Here, the UC plan performance across the representative scenarios is considered. We also compare the overall UC plan performance for method 1 and method 2. The objective function for performance evaluation is chosen as:

$$f = \sum_t \sum_i g_{it} C_{it} + \gamma \sum_t (P_{L,t} - \sum_i g_{it} - g_{wt,m})^2 + \lambda \sum_t P_{loss,t})$$

Tables 7 and 8 present the performance evaluation for the two methods.

Method 1 performance evaluation

Table 7

Wind Scenario	Very High	High	Medium	Low	Very Low
Probability	0.0139	0.1314	0.4658	0.3600	0.0289
Correction Cost	116.0053	3.7191	6.8265	56.3789	338.8427
Total Cost	1319.2	1206.9	1210.0	1259.6	1542.0

Method 2 performance evaluation

Table 8

Wind Scenario	Very High	High	Medium	Low	Very Low
Probability	0.0139	0.1314	0.4658	0.3600	0.0289
Correction Cost	124.3360	5.1023	9.3974	66.3734	363.8691
Total Cost	1303.1	1183.9	1188.2	1245.2	1542.7

Analysis:

The correction cost and total cost are both much higher when the UC plan is evaluated on extreme wind generation futures. Moreover, the very low case has the largest cost. This is not surprising because, for a very low wind generation scenario, more reserve power must be used which is usually expensive. One point to notice is that, even under very high wind generation scenario, the correction cost and total cost are higher. This is because the very high wind-generation case has the lowest probability, thus, it diverges more from the expected behavior of wind generation which results a higher cost in our formulation.

Comparison: Stochastic and deterministic UC

Method 2 incurs a higher correction cost but has a lower total cost compared to Method 1. This is reasonable because method 2 use an expected cost measure which accounts the uncertainty in wind generation. Thus, the solution provided by method 2 requires more flexibility in terms of correction on the day of operations, but can reduce overall cost.

5 Conclusions

Computing technologies have been used in managing the electric power grid for many years. Yet, the growing pervasiveness of cyber-systems in the control room – which include cluster- and cloud- based systems with unprecedented computational power, increasingly high-bandwidth data communications, improved visualization technologies, etc. – can provide unique opportunities for decision-making and management, which are far from fully realized. At their essence, these new cyber- capabilities allow a seamless integration of the physical-world, human, and economic aspects of power-grid management. In this sense, they are transforming the grid from a collection of disparate processes into an integrated cyber-physical system.

In this article, we have envisioned using the growing integration of cyber-technologies in the control room, to assist in the operation of power systems with high penetration of intermittent-renewables. Specifically, we have envisioned an end-to-end framework for forecasting probabilistic renewable-generation futures, and using these futures for unit-scheduling for the day-ahead market. The envisioned framework exploits new cyber- capabilities in myriad ways: it uses ensemble forecasting data in real time to inform forecasting/design, draws on stochastic network modeling tools such as the influence model, and imagines an approach to unit-commitment that distinguishes critical units to simplify computation and aid decision-makers. Here, the envisioned framework has been prototyped for a small-scale case study, using ensemble forecast data from a historical day of interest (specifically, wind forecast data for the Columbia River Gorge area of Washington/Oregon on that day). This case study, while preliminary, indicates that the framework may lead to practical new technologies for unit scheduling on the day ahead. A crucial next step is to compare the performance of the proposed methods with benchmark methods for a larger-scale example.

References

1. Ott, A.L.: Experience with PJM Market Operation, System Design, and Implementation. IEEE Transactions on Power Systems 18(2) (May 2003)
2. CAISO Smart Grid Roadmap and Architecture, Publication of the California Independent System Operator (CAISO) (December 2010)
3. NYISO Transmission and Dispatching Operations Manual (Manual 12), Publication of the New York Independent System Operator (October 2012)
4. PJM Manual 11: Energy and Ancillary Services Market Operations, Prepared by Forward Market Operations group at PJM (2013)
5. Borenstein, S.: The trouble with electricity markets (and some solutions). Working paper of the Program on Workable Energy Regulation, University of California Energy Institute (January 2001)
6. Hawkins, D., Rothleder, M.: Evolving role of wind forecasting in market operation at the CAISO. In: Proceedings of the IEEE Power Systems Conference and Exposition (October 2006)
7. Xie, L., Carvalho, P., Ferreira, L., Liu, J., Krogh, B., Popli, N., Ilic, M.: Wind integration in power systems: operational challenges and possible solutions. Proceedings of the IEEE 99(1), 214–232 (2011)

8. Smith, J.C., Milligan, M.R., DeMeo, E.A., Parsons, B.: Utility wind integration and operating impact state of the art. IEEE Transactions on Power Systems 22(3), 900–908 (2007)
9. Lange, M.: On the uncertainty of wind-power predictions – analysis of the forecast accuracy and statistical distribution of errors. Transactions of the ASME-N-Journal of Solar Energy (June 2004)
10. Wu, Y.-K., Hong, J.-S.: A literature review of wind forecasting technology in the world. In: Proceedings of IEEE PowerTech, Lausanne, Switzerland (July 2007)
11. Tuohy, A., Meibom, P., Denny, E., O'Malley, M.: Unit commitment for systems with significant wind penetrations. IEEE Transactions on Power Systems 24(2), 592–601 (2009)
12. Botterud, A., Wang, J., Monteiro, C., Miranda, V.: Wind power forecasting and electricity market operations. Proceedings of USAEE 3, 3846 (2009)
13. Orwig, K.D., et al.: Enhanced short term wind power forecasting and value to grid operations. In: Proceedings of the 11th Annual International Workshop on Large Scale Integration of Wind Power into Power Systems, Lisbon, Portugal, November 13-15 (2012)
14. Jiang, J., Roy, S.: Stochastic prediction of spatio-temporal solar-generation futures: an influence-model-based methodology, `http://www.eecs.wsu.edu/~sroy` (in preparation)
15. Asavathiratham, C., Roy, S., Verghese, G.C., Lesieutre, B.C.: The influence model. IEEE Control Systems Magazine (December 2001)
16. Roy, S., Wan, Y., Taylor, C., Wanke, C.R.: A Stochastic Network Model for Uncertain Spatiotemporal Weather Impact at the Strategic Time Horizon. In: Proceedings of 10th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference, Fort Worth, TX (September 2010)
17. Xue, M., Zobell, S.M., Roy, S., Taylor, C., Wan, Y., Wanke, C.: Using stochastic, dynamic weather impact models in strategic traffic flow management. In: Proceedings of the Second Aviation, Range and Aerospace Meteorology Special Symposium on Weather-Air Traffic Management Integration, Seattle, WA (January 2011)
18. `http://nomads.ncep.noaa.gov/txt_descriptions/SREF_doc.shtml`
19. Wan, Y., Roy, S., Lesieutre, B.C.: Uncertainty evaluation through mapping identification in intensive dynamic simulations. IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans 40(5), 1094–1104 (2010)
20. Xue, M., Roy, S., Zobell, S.M., Wan, Y., Taylor, C., Wanke, C.: A Stochastic Spatiotemporal Weather-Impact Simulator: Representative Scenario Selection. In: Proceedings of the 2011 Aircraft Technology Integration and Operations Conference, Virginia Beach, VA (September 2011)
21. Barth, R., Brand, H., Meibom, P., Weber, C.: A stochastic unit-commitment model for the evaluation of the impacts of integration of large amounts of intermittent power. In: Proceedings of the 9th International Conference on Probabilistic Methods Applied to Power Systems, Stockholm, Sweden, June 11-15 (2006)
22. Takriti, S., Krasenbrink, B., Wu, L.S.-Y.: Incorporating Fuel Constraints and Electricity Spot Prices into the Stochastic Unit Commitment Problem. Operations Research 48(2), 268–280 (2000)
23. Johnson, R.B., Oren, S.S.: Equity and efficiency if unit commitment in competitive electricity markets. Utilities Policy 6(1), 9–19 (1997)
24. ERCOT – Generation, `http://www.ercot.com/gridinfo/generation/`

25. Cook, S.R., Gelman, A., Rubin, D.B.: Validation of software for Bayesian models using posterior quantiles. Journal of Computational and Graphical Statistics 15(3), 675–692 (2006)
26. Wu, T., Rothleder, M., Alaywan, Z., Papalexopoulos, A.D.: Pricing energy and ancillary services in integrated market systems by an optimal power flow. IEEE Transactions on Power Systems 19(1), 339–347 (2004)
27. Lesieutre, B.C., Oh, H., Thomas, R.J., Donde, V.: Identification of market power in large-scale electric energy markets. In: Proceedings of the 39th Hawaii International Conference on Systems Science (January 2006)
28. Meibom, P.: Stochastic Optimization Model to Study the Operational Impacts of High Wind Penetrations in Ireland. IEEE Transaction on Power Systems 26(3) (August 2011)
29. Morgan, E.C.: Probability distributions for offshore wind speeds. Energy Conversion and Management 52 (2011)
30. Wood, A.J.: Power Generation, Operation and Control. Wiley (1996)
31. Soliman, S.A.: Modern Optimization Techniques with Applications in Electric Power Systems. Springer (2011)
32. McCalley, J.D.: Lecture on Unit Commitment. Personal Collection of EE553, Iowa State University, IA (2012)
33. Drgrib (NDFD GRIB2 Decoder), `http://www.nws.noaa.gov/mdl/degrib/txtview.php?file=tkdegrib.txt&dir=base`
34. Hodge, B.: Wind Power Forecasting Error Distributions over Multiple Timescales (2011), `http://www.nrel.gov/docs/fy11osti/50614.pdf` (retrieved)
35. Generation (2014), `http://www.ercot.com/gridinfo/generation/` (retrieved)
36. European Centre for Medium-Range Weather Forecasts, `http://data-portal.ecmwf.int/data/d/interim_daily/` (retrieved)
37. Liggett, T.M.: Interacting Particle Systems. Springer, New York (1985)
38. Constantinecu, E.M., Zavala, V.M., Rocklin, M., Lee, S., Anitescu, M.: A computational framework for uncertainty quantification and stochastic optimization in unit commitment with wind power generation. IEEE Transactions on Power Systems 26(1), 431–441 (2011)
39. Ruiz, P.A., Philbrick, C.R., Zak, E., Cheung, K.W., Sauer, P.W.: Uncertainty management in the unit commitment problem. IEEE Transactions on Power Systems 24(2), 642–651 (2009)
40. Basu, S., Choudhury, T., Clarkson, B., Pentland, A.: Learning human interactions with the influence model. In: Neural Information Processing Systems (2001)
41. Khaitan, S.K., McCalley, J.D.: Cyber physical system approach for design of power grids: A survey. In: 2013 IEEE Power and Energy Society (PES). IEEE (2013)
42. Khaitan, S.K., McCalley, J.D.: Design techniques and applications of cyberphysical systems: a survey. IEEE Systems Journal (99), 1–16 (2014)