



WIND POWER FORECASTING

Independent system operators, regional transmission organizations, and utilities all have a critical need for accurate forecasts of electrical generation by wind.

By Daniel Kirk-Davidoff

Daniel Kirk-Davidoff is chief scientist for weather and climate services, MDA Information Systems, Inc., and adjunct assistant professor, Department of Atmospheric and Oceanic Science, University of Maryland, College Park. Go online to www.mdainformationsystems.com.


ELECTRICAL GENERATION BY WIND TURBINES

has grown dramatically over the past decade. Construction of wind turbines has been an attractive option for developers because the turbines generate no onsite pollution and thus receive preferential treatment of various forms from governments at various levels, they require no substantial inputs once installed, and they are constructed in small units of tens of megawatts of generating capacity, requiring much smaller upfront costs than typically sized fossil fuel or nuclear power plants. As a result, in the Electrical Reliability Council of Texas (ERCOT) service area, electricity from wind tur-

bines now frequently exceeds 20 percent of total electrical demand during the winter, when winds are strong and electrical demand relatively low.

WIND VARIABILITY

The challenge to electrical system operators from wind comes from its variability. Wind generation variability now represents a large share of the total variability of the net electrical demand (demand minus renewable generation) that must be met by non-renewable power sources. Neither demand variability nor renewable generation is perfectly predictable, but it is safe to say that wind genera-



tion represents a major source of uncertainty in the net demand forecast in either the next hour or the next day. For example, in ERCOT in the winter, wind generation typically fluctuates between 1 to 7 GW over the course of a day or so, while total demand varies between 24 to 36 GW over the same period.

Independent system operators, regional transmission organizations, and utilities all have a critical need for accurate forecasts of electrical generation by wind. Each increment in forecast accuracy allows a reduction in the need for costly last-minute power purchases, or “curtailments,” when gen-

erators are asked, and sometimes paid, to cut back generation.

Wind generation forecasts are especially in demand for two discrete lead-time intervals. The day-ahead forecast—issued at least once on the morning of one day, and valid at hourly intervals for the period 12:00 a.m. through 11:59 p.m. of the following day—is used to set conditions for bidding on delivery of electrical generation in the many electrical markets that auction off the right to sell power into the grid for each hour of the following day. The short-term forecast—issued at hourly or greater frequency for intervals of 15 minutes or less out to perhaps six hours in advance—is of particular interest to electrical system operators who must balance electrical supply and demand perfectly, and can save money by giving warning to suppliers of electrical generation in the “spot” market for last minute generation requests. These forecast customers are particularly interested in accurate forecasts of wind “ramps,” when wind generation increases or decreases rapidly over a period of minutes to hours. Forecasts of wind generation are in demand by electrical system operators (utilities, regional transmission organizations, and independent system operators), by participants in electrical markets (electrical generation owners, traders in natural gas, and other fuels), and by wind farm owners and operators, who may be required provide forecasts of their expected generation, or may need forecasts in order to offer bids in the day-ahead market.

At MDA Information Systems, we provide forecasts of wind generation for individual wind farms and for aggregate generation over all wind farms in a Regional Transmission Organization territory (currently for MISO, ERCOT, CAISO, PJM, BPA, AESO, and IESO). Forecasts are based on a continuously tuned ensemble of forecast models, and on statistical prediction based on observed winds from the wind farm and the neighboring observation stations.

FORECASTING TECHNOLOGY

The ability to forecast the wind speed and direction at all locations in the atmosphere is at the heart of modern weather forecasting. Only because we can accurately model the forces that operate on the atmosphere, and predict the resulting changes in momentum of air at each location, are we able to predict where upward motion will generate cooling, condensation, clouds, and rain; where downward motion will produce warming, drying, and sunny weather; where winds from the tropics will tend to produce a warm day; and winds from the polar regions a cold day. Our ability to predict the track and intensity of winter storms before they even come into existence depends entirely on the

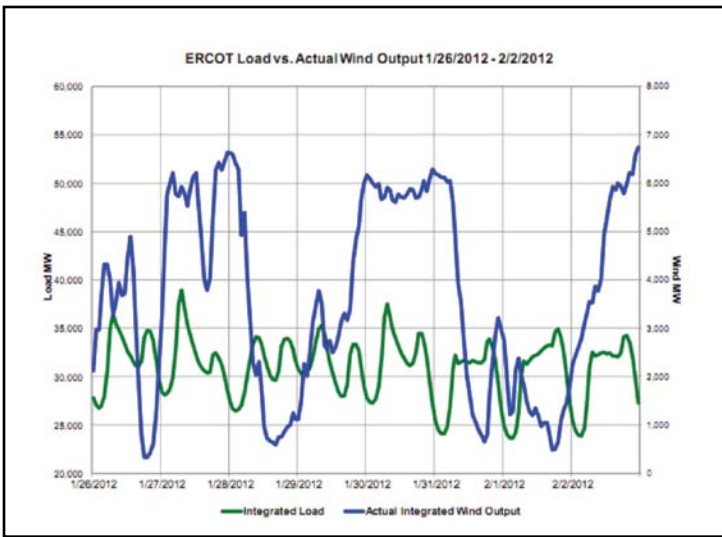


Fig. 1: Wind generation (scale is five times larger) and total electrical load in ERCOT over a week in late January 2012. (Courtesy of ERCOT)

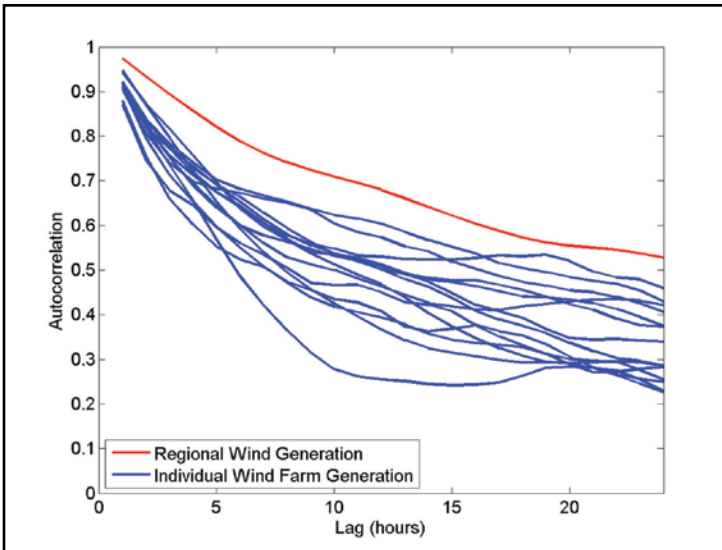


Fig. 2: Lag-autocorrelation of wind generation at individual wind farms and of aggregate wind generation for all farms, for the AESO interconnection. Data for the month of January, 2012 (AESO).

existence of computer models of the atmosphere that integrate Newton’s equations of motion at high resolution around the world and throughout the depth of the atmosphere.

Without these Numerical Weather Prediction (NWP) methods, forecasting must rely on statistics. For example, in many locations wind speeds at

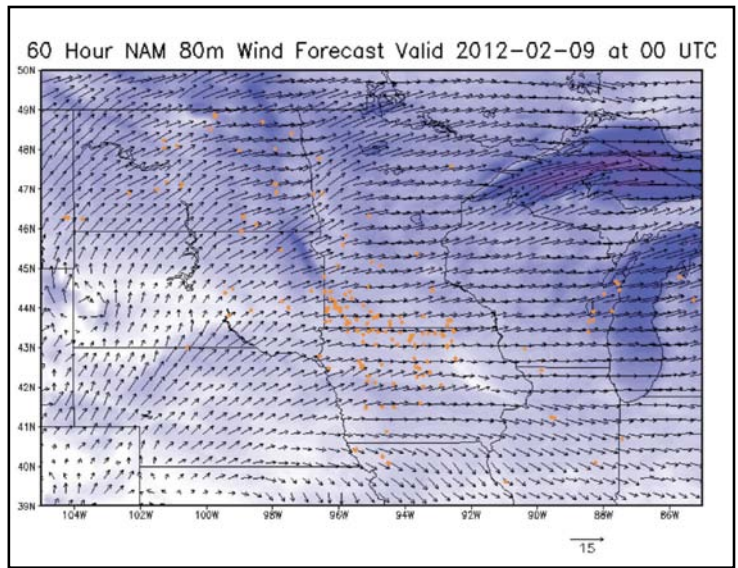
the height of a wind turbine hub (about 80 m above the ground) are stronger at night than in the day, and stronger in the winter than in the summer. In addition, like variations in most geophysical variables, variations in wind have some “memory.” The wind speed 20 minutes from now is much more likely to be about as strong as it is right now than it

is to be much weaker or much stronger. Thus, a forecast of the wind generation one hour from now might be made based only on lag-regression methods that detect relationships between the wind generation at a given time with the wind generation at a few time intervals before that given time. The lag regression for individual wind farms is lower than for aggregate generation over a large number of farms (fig. 2), which makes forecasting for individual farms more difficult.

For very long-range forecasts and very short-range forecasts, we must rely at least partially on statistics. Because of the chaotic nature of the atmosphere, the detailed motions are unpredictable beyond about two weeks: the atmosphere does not yet “know” if it will be a calm day or a windy day two weeks from now at any given location. However, computer models can make use of information about the state of the oceans, which have a long “memory,” to predict that the weather a couple of weeks or months from now is likely to be windier or less windy than average. Even without NWP models we can make very long-range forecasts based on persistent climate features like the El Niño/La Niña variation. If the Pacific Ocean is currently in a strong La Niña state, it is likely to still be in a strong La Niña state two months from now. If the winds in a particular location tend to be stronger under La Niña than under average conditions, then we can skillfully predict that winds a few months from now are likely to be stronger than average in that particular location.

For the short-range forecast, NWP methods must compete with the atmosphere’s strong memory. Since a prediction that the wind in the next hour will be the same as the wind right now is likely to be a very good forecast, NWP methods must be very powerful to improve upon

Fig. 3: Wind speed (indicated by color saturation, and wind barb length) and direction (arrows point in the direction of wind flow) derived from the U.S. National Center for Environmental Prediction North America Model (NAM). The wind farms shown in orange are those that contribute to the Midwest Independent System Operator (MISO) electrical interconnection.



statistical methods. As computers increase in speed, it becomes possible to incorporate more and more data (e.g. from Doppler radar) into the NWP models more and more quickly, and to run the models at ever-higher resolution. This allows NWP methods to compete with statistical methods at shorter and shorter lead times, offering the hope that a rapid increase in wind speed at one wind farm—due, for example, to a gust of wind emerging from a single strong thunderstorm, and causing a “ramp” in wind electrical generation—might be predictable a few hours in advance.

At present, short-term wind power forecasts typically use statistical techniques that learn from the data on wind speed and power variations at the location of interest and surrounding nearby observation points, and develop predictive relationships between present and recent winds at and around this location and future winds. This allows more skillful forecasts of the wind than NWP methods for up to three hours. In some specialized situations, where the short range forecast is of particular interest, and where complex terrain makes the atmospheric flow field complex, very high resolution NWP mod-

els can be run over a small region (perhaps a square region a few hundred kilometers on a side), and can contribute skill even on these short time scales.

Beyond that time, we use a collection (or ensemble) of NWP models to forecast wind speed from three hours out to 10 days or more. The NWP models predict wind, temperature, water vapor, and precipitation at distinct locations on a grid that covers a large area—either the whole world, or some discrete region on the earth. To predict wind generation from an individual wind farm, each model’s wind speed prediction is first interpolated from the model grid to the location of the wind farm. Figure 3 shows the wind speed and direction prediction for a single NWP model, the U.S. National Weather Service North America Model, with the locations of wind farms shown as orange dots. The type of wind turbine at the farm is then noted, and that turbine’s power curve, the relationship between wind speed at the turbine hub and the wind power generated by the turbine, is used to estimate the power produced by all the turbines in the farm. If the farm extends over more than a single model grid square,

this operation may be repeated for each individual turbine, and the production from all turbines summed to predict the wind farm generation.

By comparing each NWP model’s prediction with observations over the past several weeks, we remove biases from each model. Combining these tuned forecasts, and weighting them by their accuracy, so that more skillful models receive more weight than less skillful models, produces a consensus forecast more skillful than any of the individual models. The range of predictions of the full set of tuned models can be used to estimate the uncertainty of that particular forecast. When the full set of models predicts a wide range of wind speeds at a given time in the future, our confidence in the prediction of the model consensus should be somewhat lower than when all the model predictions fall tightly together.

Our forecast processes are summarized in fig. 4, which shows the flow of data from the client (who provides actual wind and wind power generation data at the wind farm, in real time if possible), to our company, and back to the client as a wind and wind energy forecast. The key to



Fig. 4: Flow chart of MDA's wind power forecast system.

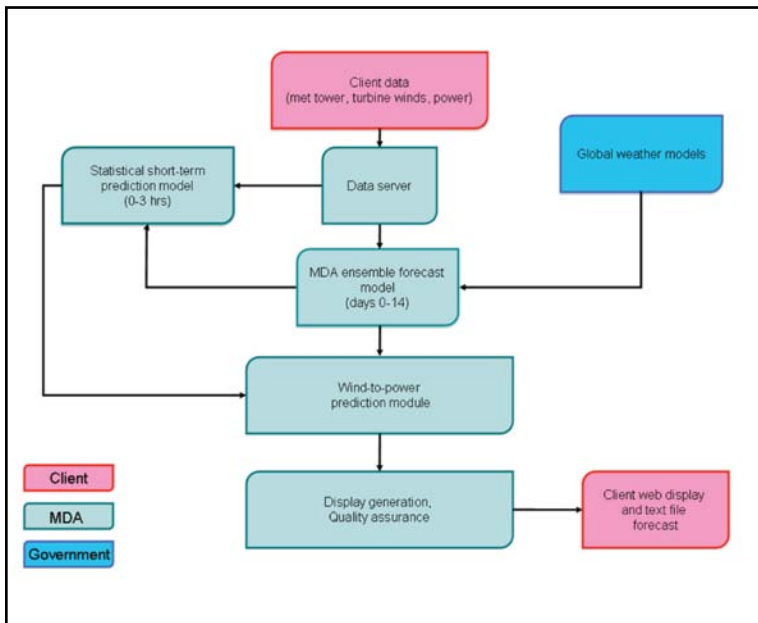


Fig. 5: MDA's wind generation forecast Web interface.

skillful wind power forecasting is the care and research that is put into the process of bias removal from each model.

This ensemble forecasting methodology also allows skillful prediction of forecast uncertainty. The magnitude of past forecast errors are examined by lead time, and correlated with the range of forecast from our tuned model ensemble. Once derived, this relationship between model spread and forecast errors in the past allows us to skillfully predict the uncertainty of each forecast at each lead time. Detailed

knowledge of forecast uncertainty allows energy managers to minimize their financial risk from electricity markets.

Wind forecasts are presented to users as a Web display, from which numerical data can be generated for incorporation into spreadsheets or other software. An example is shown in fig. 5. The interface allows a user to view the multiple tuned model forecasts that underlie the consensus forecast, and also shows confidence intervals (\pm one standard deviation). The menu allows the user to select the region or wind

farm of interest, while buttons allow the user to check past forecast history and skill.

FUTURE PROGRESS

To date, much wind generation forecast skill has derived from the relatively low resolution global models run by the various national and multinational weather services, in particular the U.S. Global Forecast System and the European Center for Medium Range Weather Forecasting global model. Regional models, such as the U.S. RUC and NAM models may have lower biases, but their skill (the correlation of the forecast wind with the observed wind at a given lead time) has not been markedly better. However, as greater and greater computational and observational resources are put into improving the short-term forecast, new generations of computer models such as the High Resolution Rapid Refresh (HRRR) model, are showing good results. The HRRR issues forecasts each hour for each 15-minute interval out to 15 hours in the future. A public-private effort funded by the U.S. Department of Energy and the National Oceanographic and Atmospheric Administration, the Wind Forecast Improvement Project (www.esrl.noaa.gov/psd/psd3/wfip) is testing the utility of a significant expansion of the observational network for winds in the lower atmosphere in producing improved forecasts, with lower mean errors and improved prediction of forecast skill. ✎