

Predicting Wind Power with Greater Accuracy

Researchers are combining fieldwork, advanced simulation, and statistical analysis to help wind farm and electric power grid operators.



OVER the past decade, towering farms of wind turbines, some taller than a 40-story building, have become a fixture of the American countryside. These machines, which cleanly extract kinetic energy from the wind and convert it to electricity, today provide about 4 percent of the total electricity generated in the U.S.

President Barack Obama has established a goal of generating 20 percent of the nation's electricity from wind energy by 2030. "We believe it is reasonable to achieve that goal, because of the current high rate of wind turbine deployment nationwide," says Livermore mechanical engineer Wayne Miller, associate program leader for wind and solar power. Between 2008 and 2012, wind power capacity grew by 167 percent. "The market for new wind energy

production is complex and determined by many factors, including federal tax credits for renewable energy production and the availability of cheap natural gas," says Miller. "However, wind is now very competitive with all other sources of power generation that have been recently installed."

Miller notes that the wind is a varying and uncertain power source, dependent on a host of complex atmospheric forces. Reducing the uncertainty in wind power forecasts, on which wind farm operators and electric power grid operators rely, is the goal of a team of Livermore researchers. Many wind farms generate less energy than expected because of uncertainties in forecasting winds and in simulating the complex flows within the turbine farms. Greater understanding

of the wind is needed to optimize power production from these farms and to improve the fidelity of forecasting models that relate power output to atmospheric conditions. A major focus for the team is to better understand how power production is related to atmospheric variables, such as wind speed and turbulence, across a broad range of spatial and temporal scales and in widely varying geographic areas.

The Laboratory's wind-forecasting effort involves about a dozen atmospheric scientists, mechanical and computational engineers, and statisticians who combine fieldwork, advanced simulation, and statistical analysis. "It's a big team effort," says Miller. Partners include the National Renewable Energy Laboratory, National Center for Atmospheric Research,



University of Colorado at Boulder, Sandia and Pacific Northwest national laboratories, University of Wyoming, University of Oklahoma, University of California at Berkeley, U.S. Army, and the wind power industry. Funding is provided by the Department of Energy's (DOE's) Office of Energy Efficiency and Renewable Energy, Livermore's Laboratory Directed Research and Development (LDRD) Program, and industrial partnerships.

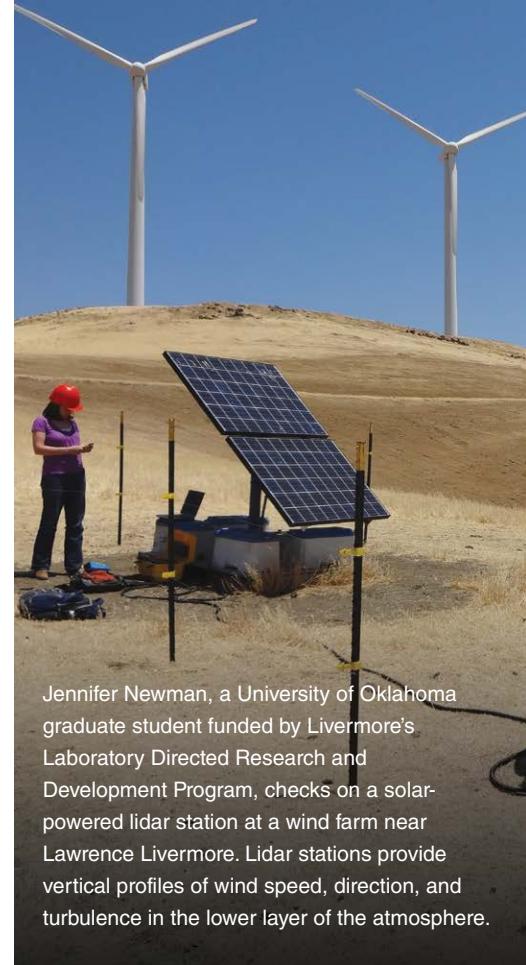
The team's advanced numerical models, verified by fieldwork and statistical analysis, account for atmospheric complexities both horizontally across landmasses and vertically above Earth's surface. High-resolution computer simulations provide useful data to wind farm developers and operators, enabling them to better select wind farm locations and operate the sites with increased efficiency. Improving the accuracy of wind predictions is also critical to electric grid operators who must dynamically balance the variable power generated by increasing or decreasing power production from other sources such as coal, natural gas, hydroelectric, geothermal, and biomass.

Today, wind power provides about 19 gigawatts, or about 4 percent, of the total electricity generated in the U.S. The Obama administration has established the goal of generating 20 percent of the nation's electricity from wind energy by 2030.

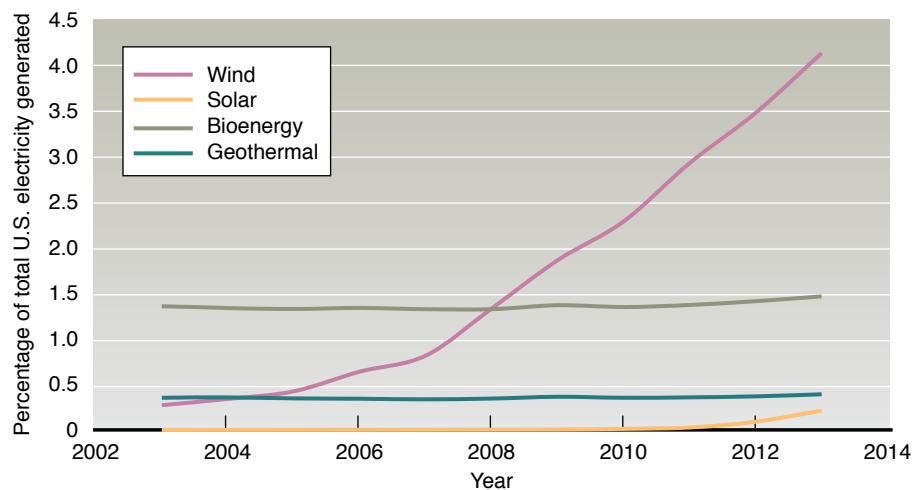
Observations Feed, Verify Simulations

Livermore field researchers are characterizing winds in numerous locations across the U.S. and are especially focused on understanding the flow in the Altamont Hills just east of Livermore, California—an area populated with wind farms. They have made significant findings studying the dynamics of atmospheric instability and turbulence. Turbulence is a particularly important variable because it affects the amount of power extracted from wind turbines as well as the reliability and life span of turbine components.

The Livermore researchers compile wind data from stationary towers mounted with weather instruments. They also use high-resolution, remote-sensing instruments including sodar and lidar to provide vertical profiles of wind speed, direction, and turbulence in the lower layer of the atmosphere. These data are collected at numerous sites such as Lawrence Livermore's Site 300 experimental research facility located about 20 kilometers east of the Laboratory's main site, DOE's Southern Great Plains Atmospheric Radiation Measurement Climate Research Facility in Oklahoma, and wind farms in the Altamont Hills and northern Oklahoma.



Jennifer Newman, a University of Oklahoma graduate student funded by Livermore's Laboratory Directed Research and Development Program, checks on a solar-powered lidar station at a wind farm near Lawrence Livermore. Lidar stations provide vertical profiles of wind speed, direction, and turbulence in the lower layer of the atmosphere.



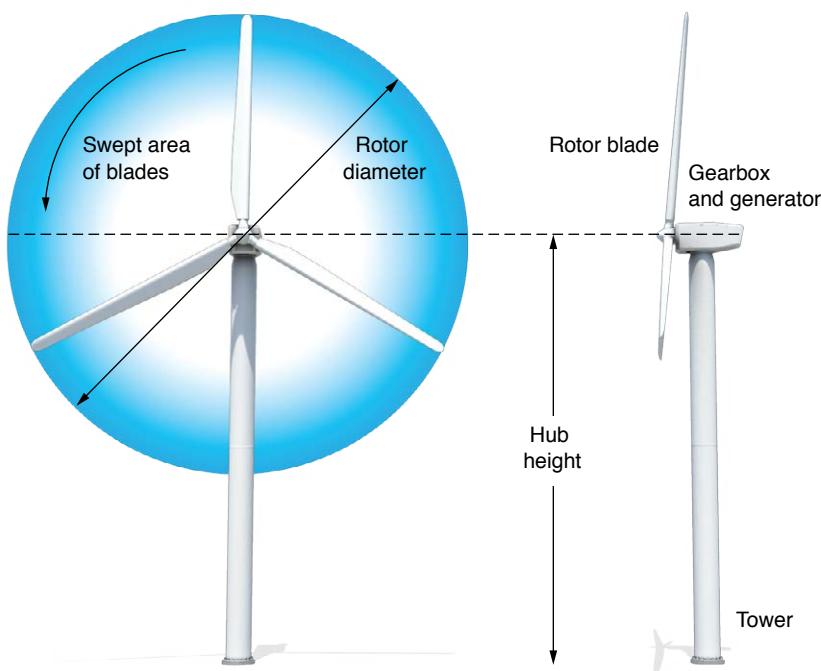


Atmospheric scientist Sonia Wharton explains that wind turbines operate in the first 150 meters of the 1-kilometer-high atmospheric boundary layer. In this layer adjacent to the ground surface, a significant exchange of heat occurs between the surface and atmosphere during daytime, which induces turbulence. Friction from the wind moving over hills, trees, and buildings also induces turbulence.

Wharton notes that as deployment of wind farms has increased, so have the average turbine hub height (distance from the ground to the blade rotor), blade diameter, and power-generating capacity. Taller turbines typically encounter higher wind speeds, allowing them to extract greater amounts of energy. Tall turbines also experience more complex airflow patterns driven by turbulent mixing. Although the average wind speed in the blade-swept area largely determines the amount of power generated (power is proportional to the cube of the wind speed), wind shear and turbulence also affect power output. For example, wind speeds can vary significantly at opposite ends of the blades, which can cause the blades to twist and deform, reducing power output and causing premature wear.

"Our measurements help us better understand the physics of the atmospheric boundary layer," says Wharton. "Increased understanding can help optimize power generation from wind farms and validate our numerical models." Wharton uses wind profile data to investigate stability factors, including vertical and horizontal turbulence intensity, veer (change in direction with height), and shear (change in wind speed with height). She compares those data to supervisory control and data acquisition (SCADA) information remotely transmitted from turbines. SCADA data, typically generated at 10-minute intervals, provide many turbine performance factors, including the yaw angle as the nacelle points to the wind, the blade pitch angle (which controls rotor speed and torque), hub-height wind speed, and power. Her research, in collaboration with scientists at the University of Oklahoma and the University of Wyoming, involves some of the most detailed studies to date exploring the relationship between three-dimensional turbulence and turbine power production.

Field studies at Site 300 and at an Altamont Hills wind farm focus on



As wind turbine blades rotate, they spin a shaft, which is connected to an electric generator. Wind turbines include a gearbox to increase the spinning shaft's speed from about 30 to 60 revolutions per minute (rpm) to the 1,000 to 1,800 rpm required by most generators to produce electricity. The hub height of most land-based turbines ranges from 60 to 100 meters. Rotor diameters can exceed 120 meters. (Rendering by Kwei-Yu Chu.)

understanding the complex wind patterns occurring in a hilly, coast-influenced region—a location similar to that of many California wind farms. Scientists have also analyzed wind profiles at flat-terrain sites in the nation, such as the very windy Great Plains, to study low-level jets and other drivers of complex atmospheric flow. “These jets, which are similar to a river of very fast air, occur at night in the Plains states and provide more power than higher level jets,” says Wharton. With improved characterization of the interactions between particular inflows and turbines over a variety of terrains, scientists and engineers can better model and design turbines for optimal performance.

Along with Laboratory atmospheric scientist Matthew Simpson and industry collaborators, Wharton is also using lidar data to study how power generation is affected by a wind ramp, which is a significant change in wind speed over just a few hours. “Our work underscores the benefit of observing complete profiles of wind speed and turbulence across the turbine rotor disk, which is available only with remote-sensing technologies such as sodar and lidar,” says Wharton. She notes that wind turbine manufacturers typically provide operators with a “power curve,” a graph that shows power from the turbine as primarily the cube of hub-height wind speed alone.

“Power curves oversimplify reality,” says Wharton. “In fact, they frequently err by plus or minus 20 percent of actual power output. We’re adding refinements to power curve models so that they reflect our knowledge of the aerodynamic environment. The more we understand atmospheric processes, the more accurate our power predictions.”

Simulations Face Challenges

Livermore’s wind-forecasting simulation and modeling efforts rely on massively parallel supercomputers to study atmospheric flows relevant to wind farm

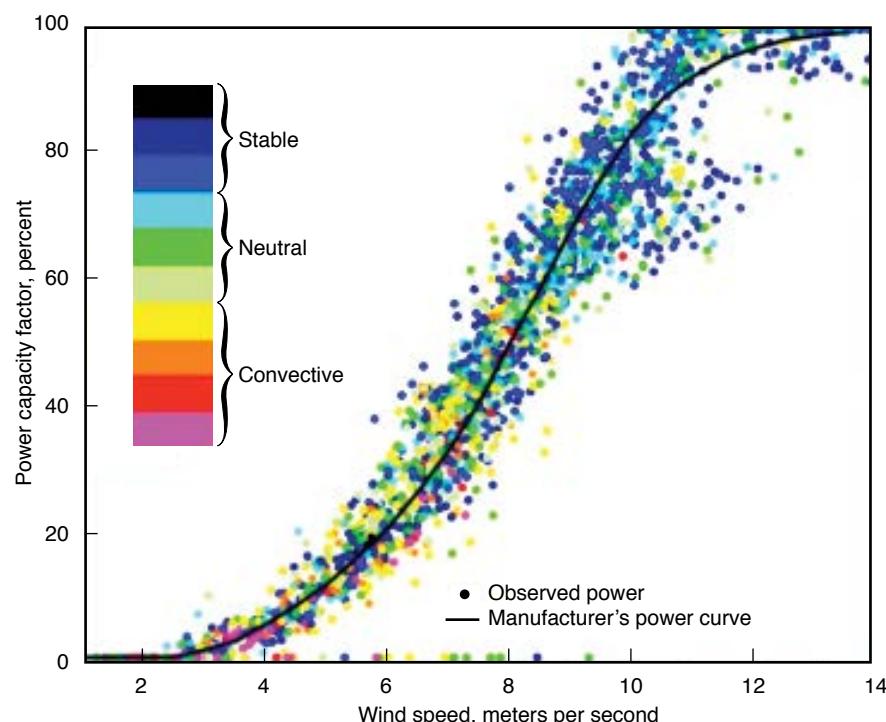
operations. The task is enormous because the length scales involved can span eight orders of magnitude—from the mesoscale (about 100 kilometers) to wind farm scales (1 to several kilometers) to turbine blade aerodynamic features (millimeters to meters).

“Simulating wind change and its effects on turbines is challenging because of the complex forces driving wind,” says Miller. “We’re essentially simulating a fluid flow in an environment where factors such as aerosols, clouds, humidity, surface-atmosphere energy exchange, and terrain influence to varying degrees both the complexity of the flow and how much power can be extracted by a spinning turbine.” Turbine rotors spin via aerodynamic forces (lift and drag) imparted to the blades when air flows

over them. The blades are attached to a shaft that, in turn, spins a generator to produce typically between 1.5 and 3 megawatts of power at full speed in land-based turbines.

Modeling the atmospheric forces that drive the wind as it acts on a single turbine poses a significant computational challenge. The job is vastly compounded when attempting to model a wind farm with 100 or more turbines. What’s more, simulations must account for varying terrain that can significantly affect power output from one wind turbine to the next. They must also account for turbulent wakes from the front rank of spinning turbine blades, which can rob power from turbines downstream.

“Because of the complexity of wind patterns and the huge range of



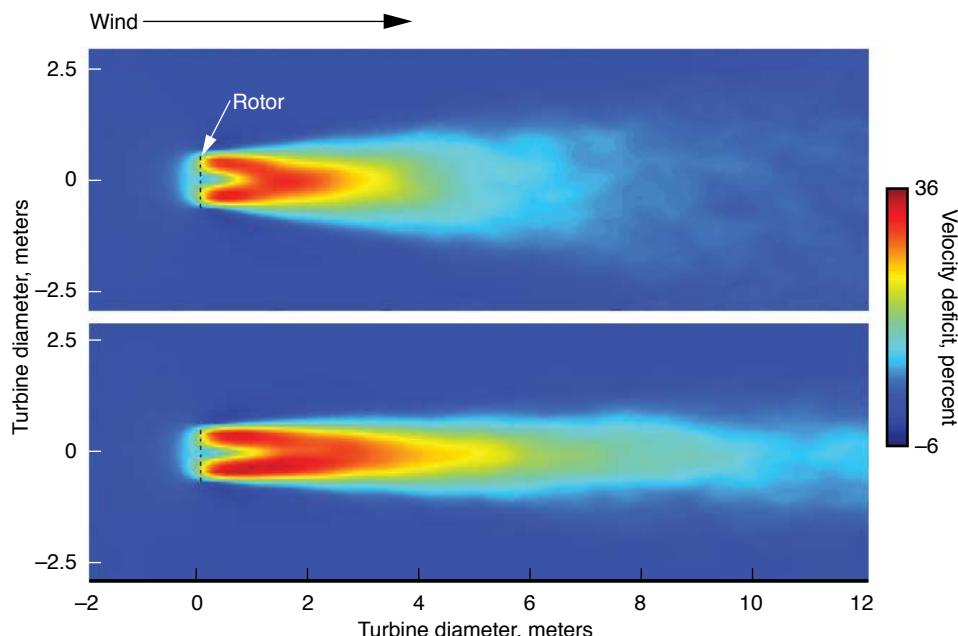
Wind turbine manufacturers typically provide operators with a simple “power curve,” which shows power from the turbine as primarily the cube of hub-height wind speed. However, Livermore researchers are showing that power curves frequently err by ± 20 percent of actual power output, as seen in this plot of observed power versus wind speed at a northern California wind farm. The color map relates atmospheric stability conditions to reported power-output observations.

relevant scales, the computational task of simulating a wind farm is daunting,” Miller says. To address knowledge gaps and research challenges associated with these simulations, atmospheric scientist Jeff Mirocha and others are extending the applicability of the Weather Research Forecasting (WRF) modeling system to the wind farm scale. This popular model, maintained collectively by more than 10,000 users and contributors worldwide, was developed primarily for larger-scale weather applications.

Mirocha says, “We have modified WRF extensively to make it applicable to the smaller scales. Accurate wind power forecasts often require a multiscale simulation approach to account for all the important scales.” As an example of multiscale methodology, one can begin with a simulation of the western U.S. to capture the evolution of large-scale weather patterns. Thereafter, a combination of smaller grid spacing and Livermore-developed submodels can accurately resolve the additional smaller-scale features that affect wind farms.

Once smaller scales of flow are resolved, wind turbine models can be implemented to investigate processes important to engineering applications such as wakes, power production, and turbine component fatigue. “Wind power simulation lies at the boundary of engineering and atmospheric sciences,” says Mirocha, who is attempting to seamlessly blend WRF atmospheric simulation with scales of motion traditionally handled by computational fluid dynamics (CFD) codes. Typical CFD codes do not contain representations of many important atmospheric physical processes contained within WRF. The complex interactions of these physical processes can strongly modulate the flow entering a wind farm and therefore the power produced or fatigue experienced by turbines.

To capture the interactions between wind turbines and complex atmospheric



Two simulations of a Generalized Actuator Disk wind turbine model within the Weather Research and Forecasting (WRF) code depict the wake downstream from a wind turbine rotor plane (dashed line). The simulations, viewed from above looking down at the turbines, show how the strength of background atmospheric turbulence (convective instability) influences turbine wakes. Stronger turbulence (top) attenuates the wake more rapidly than weaker turbulence (bottom).

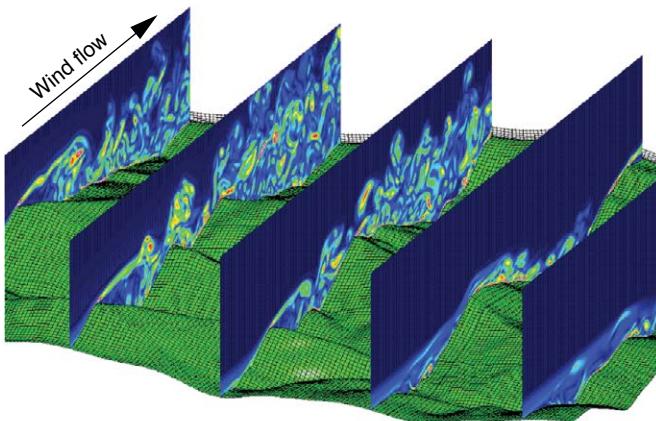
flows, Mirocha implemented a Generalized Actuator Disk (GAD) wind turbine model within WRF. GAD, an approach typically used in engineering CFD codes, depicts a two-dimensional disk containing a rotating turbine. Lift and drag forces on the turbine blades respond to atmospheric flow, including the effects of large eddies (swirling vortices of air). GAD calculates the power output of a front-rank turbine rotor and the effects from wakes that emanate downstream. These wakes, which feature both reduced wind speeds and increased turbulence, are of key concern because they are associated with power losses of up to 40 percent, and they shorten the operational life span of turbine components. WRF–GAD is being extended to a model that resolves each turbine blade independently, which will provide higher fidelity representations of blade–airflow interactions. In addition, a structural response model is being implemented within WRF to investigate

turbine fatigue in response to wakes and atmospheric turbulence.

The downscaling capability in combination with these engineering submodels will allow researchers to study a variety of phenomena unique to the wind farm environment. Says Mirocha, “The simulation framework we are developing will provide advanced tools to address these knowledge gaps, leading to improved operations, longer component life spans, and ultimately cheaper electricity.”

In addition to atmospheric physics, complex terrain can also influence the flow and turbulence experienced by wind turbines. “Eddies in the atmosphere spin and twist similar to water passing over rocks in a fast-moving stream,” explains computational scientist Kyle Chand. “A wind farm design must account for how wind is influenced as it flows past a particular terrain.” Because the standard WRF model was designed primarily for larger scales, it was restricted to simple

A simulation using Livermore's CGWind code shows the evolution of large eddies as they propagate over complex terrain at the Laboratory's Site 300 experimental research facility. The wind flows in the direction of the simulated slices, from lower left to upper right. The colors denote the strength of the turbulence, with red representing the strongest.



terrain with shallow slopes. However, an additional Livermore-developed approach, the immersed boundary method (IBM), eliminates this restriction.

The researchers combined WRF and IBM to simulate wind over hills and other complex terrains such as Oklahoma City. Says mechanical engineer Katie Lundquist, who developed IBM for her dissertation, "IBM allows us to use a Cartesian grid. Working with grids as small as 1 meter, we can accommodate, for the first time, simulations of flow in mountainous terrain with nearly vertical slopes without compromising accuracy." With DOE and LDRD funding, Lundquist is now refining the code.

Two additional codes, CGWind and HELIOS, are used for even smaller-scale simulations that would be impractical

with WRF. With CGWind, which Chand developed in 2009 with LDRD funding, topological features such as hills, buildings, and wind turbines can be accurately modeled. In addition, finer grids can be overlaid for areas requiring higher resolution, allowing the code to focus computational resources only where they are needed. CGWind has been used to model wind over a number of complex terrains, including the Laboratory's hilly Site 300.

HELIOS is a multiphysics rotor aerodynamics CFD code originally developed by the U.S. Army to model helicopters. The code was later adapted for wind turbines by the University of Wyoming. "Essentially, our colleagues at Wyoming turned the Army code for helicopters sideways to capture the tiny

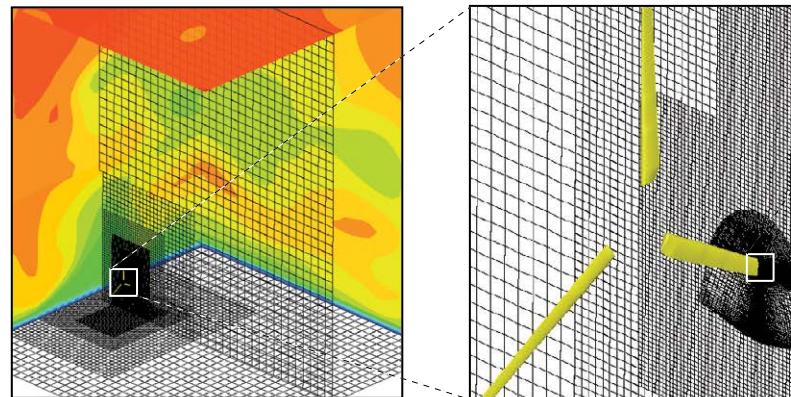
vortices created at the tips of turbine blades and to model structural loads," says Miller. In collaboration, Livermore and Wyoming have used HELIOS to perform the first-ever simulation of a 50-turbine wind farm that resolves individual spinning turbine blades. The team is extending this effort using a Livermore Computing Grand Challenge Award to run simulations of a 150-turbine farm. The simulations combine WRF and HELIOS and require up to 100,000 cores (processors working together) on the Laboratory's Vulcan supercomputer.

Minimizing Uncertainties

Another group of Livermore researchers is studying how to reduce uncertainties in the errors associated with collected field data and with the assumptions, inputs, and approximations inherent in the physics of the WRF code, its constituent modules, and the nested CFD codes. Their work takes advantage of the Laboratory's strength in statistical modeling and uncertainty quantification (UQ), which has proved invaluable in stockpile stewardship.

"Wind power forecasting involves converting atmospheric forecasts into a forecast of power output from an individual turbine or many turbines in a wind farm," says statistician Vera Bulaevskaya. Traditional manufacturer-supplied power curves model power as a function of the wind speed at the hub

Livermore and University of Wyoming simulation experts combined the WRF model with HELIOS to perform the first-ever simulation of a 50-turbine wind farm that resolves individual spinning turbine blades using realistic turbulent winds. Such realism is needed to understand how wind farms often underperform in real situations compared with idealized results. Colors show WRF wind velocity variations flowing into the HELIOS turbine grid.



height of the turbine (adjusted for air density). In reality, however, power output is a function of additional variables. For example, wind speed at heights below and above the hub, wind shear, and turbulence are also strong predictors of power production. Accounting for them provides a more complete power curve model. Bulaevskaya compares lidar atmospheric data collected at wind farms with SCADA power output data from wind turbines to gain a detailed understanding of the sensitivity of power output to changes in atmospheric conditions.

Moreover, to be valuable, a forecasting tool must not only produce accurate forecasts of power but also correctly quantify the uncertainty, or confidence level, associated with these predictions. Such confidence levels are particularly of value to electric grid operators, who need both the predictions of output and the associated levels of confidence to determine an optimal schedule for turning various sources of power on and off. Quantifying output uncertainty is also crucial for siting wind farms.

Bulaevskaya has investigated various statistical approaches for modeling power as a function of changes in atmospheric conditions. She has found that the performance of these approaches, in terms of prediction accuracy, is significantly better than that of manufacturers' power curves. (See the figure on p. 12.) One

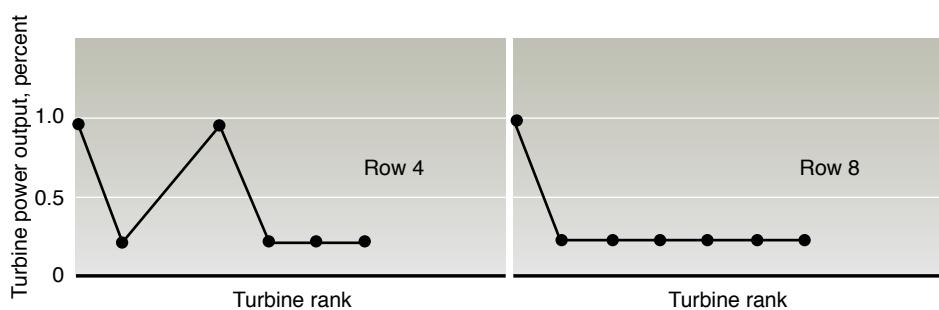
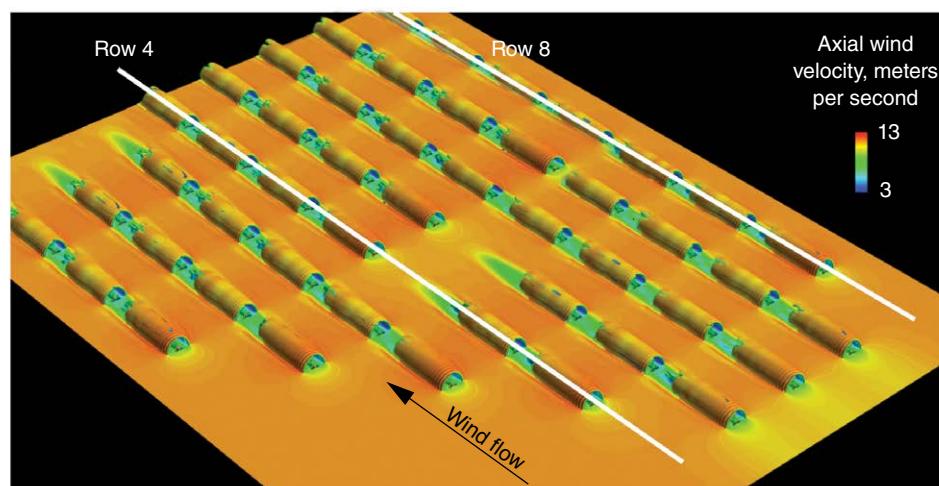
statistical technique, known as the Gaussian Process Model, is particularly well suited for estimating the uncertainty associated with predictions.

Ultimately, forecasts of atmospheric variables, rather than observed values, will be used to predict output. To reduce uncertainties in these forecasts, Simpson performs "ensemble modeling," which entails running a wind forecast model dozens or even hundreds of times using slightly different starting conditions and physics packages or submodels. (See *S&TR*, December 2013, pp. 20–23.) One ensemble can constitute 30 to 60 runs of slightly different WRF models—a process

that requires significant supercomputing resources. He explains that WRF contains many individual packages, each representing a particular atmospheric physics component, such as a model of turbulence or clouds. These packages provide descriptions of physical processes developed by different researchers.

To account for a full range of physical phenomena, Simpson varies the packages within WRF.

To model wind power predictions, Simpson begins with forecasts of atmospheric variables. The initial input is uncertain. Factors such as temperature, wind, air pressure, and large-scale weather



(top) The effect of downstream wakes caused by upstream wind turbines is illustrated in this HELIOS simulation of an offshore wind farm in Denmark. (bottom right) In row 8, turbine power is at 100 percent in the first turbine but quickly falls off because of the downstream wake. (bottom left) In row 4, the wind "recharges" in the gap after the second turbine, allowing the third turbine to produce nearly maximum power.

features ultimately affect much smaller areas such as wind farms. By running numerous simulations with a plausible sampling of varying physics packages and initial conditions, researchers obtain a range of outcomes. A set of predictions resulting from these model ensembles can then be used to quantify the uncertainty in the predictions resulting from incomplete knowledge about atmospheric physics and model inputs.

One advantage of ensembles is the capability to spot outliers such as wind ramps. Because power output is proportional to the cube of the wind speed,

a wind ramp can result in a dramatic change in power production. Consequently, accurate wind ramp prediction is extremely important, leading some experts to refer to it as “the Holy Grail of wind forecasting.” In one case study, an ensemble run showed that the ability to predict a wind ramp did not improve with grid resolution finer than 1 kilometer. Results such as these can help guide researchers to use computer resources more efficiently and further refine models.

Atmospheric scientist Don Lucas has worked extensively with climate and atmospheric model uncertainties and has

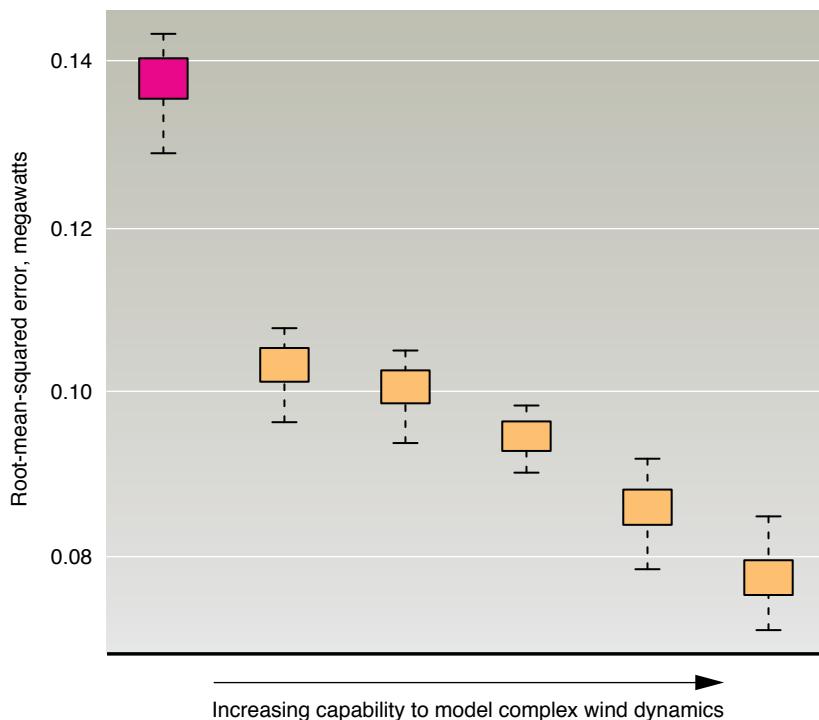
run thousands of ensemble simulations during his career. “UQ is at the interface of computer simulation and statistical analysis,” he says. “Sometimes, changing parameters or their relative strengths does not affect the output or exerts only a small influence. We can determine which factors greatly influence forecast and focus computational resources on those.” At the same time, he notes that relevant field data help keep models “honest.” “We want to improve UQ calculations with observations to see how well we know the model and how well the model performs.”

Benefits of Accurate Predictions

Together, the field observations, simulations, and statistical analyses are significantly improving wind power predictions. The Laboratory is sharing the results of this work with the wind industry to help turbine manufacturers refine their power curves and incorporate findings about what atmospheric processes are important in wind power forecasting.

With improved models, wind farm operators will know how to better maximize their sizable investments, more skillfully bid into the energy market, optimally site turbines, and minimize the turbulent wakes from upstream turbines. The biggest winner, however, may well be the American consumer, who will enjoy abundant supplies of energy from a clean and inexhaustible source.

—Arnie Heller



As statistical models account for increasingly complex information about wind dynamics, the root-mean-squared error associated with predicted power output from a wind turbine decreases significantly. Livermore researchers used lidar data as input to five power curve models (gold squares) and compared their results with a manufacturer's power curve (pink square) for a number of data sets. The power curve models incorporating successively more information about wind dynamics tended to have a much lower prediction error.

Key Words: atmospheric boundary layer, CGWind, computational fluid dynamics (CFD), electric power grid, Gaussian Process Model, Generalized Actuator Disk (GAD), HELIOS, immersed boundary method (IBM), lidar, mesoscale, sodar, supervisory control and data acquisition (SCADA), turbulence, Weather Research Forecasting (WRF), wind turbine farm, wind power forecast.

For further information contact Wayne Miller (925) 424-4472 (miller99@lbl.gov).