NREL
NATIONAL RENEWABLE ENERGY LABORATORY

A Large-Eddy Simulation of Wind-Plant Aerodynamics

Preprint

Matthew J. Churchfield, Sang Lee, and Patrick J. Moriarty
National Renewable Energy Laboratory

Luis A. Martinez and Stefano Leonardi
University of Puerto Rico

Ganesh Vijayakumar and James G. Brasseur
The Pennsylvania State University

To be presented at the 50th AIAA Aerospace Sciences Meeting
Nashville, Tennessee
January 9-12, 2012

Conference Paper
NREL/CP-5000-53554
January 2012

Contract No. DE-AC36-08GO28308

NOTICE

Available electronically at http://www.osti.gov/bridge

Available for a processing fee to U.S. Department of Energy
and its contractors, in paper, from:

U.S. Department of Energy
Office of Scientific and Technical Information

P.O. Box 62
Oak Ridge, TN 37831-0062
phone: 865.576.8401
fax: 865.576.5728
email: mailto:reports@adonis.osti.gov

Available for sale to the public, in paper, from:

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
phone: 800.553.6847
fax: 703.605.6900
email: orders@ntis.fedworld.gov
online ordering: http://www.ntis.gov/help/ordermethods.aspx

Cover Photos: (left to right) PIX 16416, PIX 17423, PIX 16560, PIX 17613, PIX 17436, PIX 17721

Printed on paper containing at least 50% wastepaper, including 10% post consumer waste.

A Large-Eddy Simulation of Wind-Plant Aerodynamics

Matthew J. Churchfield[1], Sang Lee[2], and Patrick J. Moriarty[3]
National Renewable Energy Laboratory, Golden, Colorado, 80401

Luis A. Martínez[4] and Stefano Leonardi[5]
University of Puerto Rico, Mayagüez, Mayagüez, Puerto Rico, 00681

and

Ganesh Vijayakumar[6] and James G. Brasseur[7]
The Pennsylvania State University, University Park, Pennsylvania, 16802

In this work, we present results of a large-eddy simulation of the 48 multi-megawatt turbines composing the Lillgrund wind plant. Turbulent inflow wind is created by performing an atmospheric boundary layer precursor simulation, and turbines are modeled using a rotating, variable-speed actuator line representation. The motivation for this work is that few others have done large-eddy simulations of wind plants with a substantial number of turbines, and the methods for carrying out the simulations are varied. We wish to draw upon the strengths of the existing simulations and our growing atmospheric large-eddy simulation capability to create a sound methodology for performing this type of simulation. We used the OpenFOAM CFD toolbox to create our solver.

The simulated time-averaged power production of the turbines in the plant agrees well with field observations, except with the sixth turbine and beyond in each wind-aligned. The power produced by each of those turbines is overpredicted by 25–40%. A direct comparison between simulated and field data is difficult because we simulate one wind direction with a speed and turbulence intensity characteristic of Lillgrund, but the field observations were taken over a year of varying conditions. The simulation shows the significant 60–70% decrease in the performance of the turbines behind the front row in this plant that has a spacing of 4.3 rotor diameters in this direction. The overall plant efficiency is well predicted. This work shows the importance of using local grid refinement to simultaneously capture the meter-scale details of the turbine wake and the kilometer-scale turbulent atmospheric structures. Although this work illustrates the power of large-eddy simulation in producing a time-accurate solution, it required about one million processor-hours, showing the significant cost of large-eddy simulation.

I. Introduction

USING large-eddy simulation (LES), we simulate the aerodynamics and power production of the Lillgrund wind plant,[1] which contains 48 Siemens SWT-2.3-93 wind turbines.[2] This type of turbine has a rotor diameter of 93 m and a rated power production of 2.3 MW. Because of the scale of such a computation, we only examine the effects of wind from one direction at one speed and turbulence intensity. This is not an exhaustive study, but rather is a documentation of our first attempt at simulating a full-scale, operational wind plant with comparison to actual production data.[3] Aside from exploring the scientific results concerning wake effects and power production, we also present our methodology, problems we encountered in performing the simulation, and areas for improvement in future simulations.

[1] Research Engineer, National Wind Technology Center, 1617 Cole Boulevard/MS3811, AIAA Member.
[2] Postdoctoral Researcher, National Wind Technology Center, 1617 Cole Boulevard/MS3811, AIAA Member.
[3] Senior Engineer, National Wind Technology Center, 1617 Cole Boulevard/MS3811, AIAA Member.
[4] Graduate Research Assistant, Department of Mechanical Engineering, P.O. Box 9045.
[5] Assistant Professor, Department of Mechanical Engineering, P.O. Box 9045.
[6] Graduate Research Assistant, Department of Mechanical and Nuclear Engineering, 335 Reber Building.
[7] Professor, Department of Mechanical and Nuclear Engineering, 205 Reber Building.

Wind turbines extract energy from the flow, and as a byproduct they create a wake that trails behind them. If wind turbines are arranged in an array, as with a wind plant, wakes from upwind turbines may interact with downstream turbines. Wake effects include decreased power produced by downstream turbines relative to upstream turbines. They also include increased mechanical loading due to wake turbulence and partial wake effects in which a blade operates in the wake during part of the time and in the freestream for the remainder of the time. The current state of knowledge concerning wind turbine wakes and how they interact with other turbines and the atmospheric boundary layer is not complete.

Increasing that knowledge would allow operators to more effectively run their plants and help developers to better site turbines within wind plants to increase efficiency and mitigate loads. It would enable researchers to create better turbine- and wind plant-level control systems. It would also assist manufacturers in engineering the next generation of larger turbines that will be more flexible and, unless carefully designed, susceptible to wake-induced damage. Field measurements are necessary to increase our knowledge of wind turbine wakes. Although important, experiments used to gather such data are costly and difficult given the large scale of wind plants. A lower-cost complement to these experiments is LES, the subject of this work.

LES has the potential to yield a high-resolution, both in time and space, representation of the flow through wind plants including turbine wake development and interaction with other turbines and the atmospheric boundary layer that provides the winds. It is still a research tool, though, and the best practices for applying LES to wind plant flows have yet to be fully established. The focus of this work is to present a methodology for performing LES computations of flow through a modern wind plant composed of many multi-megawatt turbines. This framework is not meant to be used as a design tool by industry, but rather as a research tool to better understand wind plant aerodynamics and wake effects so that, eventually, improved reduced-order tools needed by industry can be created.

Few have performed LES of full wind plants. One of the most prominent works is that of Ivanell.[4] He performed LES of two of the ten rows of the Horns Rev wind plant and assumed periodic conditions on the side boundaries to approximate the full plant aerodynamics. To approximate the turbines, he used advanced actuator disks with 80 m diameters that not only apply thrust force to the flow, but also tangential force. The mean inflow wind follows a power law profile, and a plane of fluctuating body forces parallel to and near the upstream boundary creates turbulence. The generated turbulence is dictated using the Mann[5] model and is meant to represent a neutrally stable atmospheric boundary layer. Ivanell uses a mesh with resolution of roughly 4 m with local 2 m resolution near the turbines. He varied the wind inflow angle by plus or minus 15° of alignment with the turbine rows, and he also varied the inflow turbulence intensity between 2% and 5%. He observed that when the simulated wind is aligned with the turbine rows, the wake effect is significantly overpredicted as compared with ten-minute averaged field data and the downstream turbine power production is underpredicted by as much as 30%. When the simulated wind is not aligned with the turbine rows, his power predictions agree better with the ten-minute averaged field data. It is unclear if the discrepancy is caused by the fact that real wind does not stay perfectly aligned with a turbine row over a period of ten minutes or if the imposed turbulence does not contain enough large-scale fluctuations to mimic these minute-scale wind direction changes.

Another prominent work is that of Calaf, Meneveau, and Meyers,[6] whose wind plant simulation is based more on the atmospheric sciences perspective. They perform a precursor atmospheric LES in which a laterally periodic domain of at least 3 km × 3 km × 1 km is used to generate a turbulent wind field of neutral atmospheric stability. The grid resolution is roughly 20 m in the horizontal directions and 7 m in the vertical. This work differs from that of Ivanell[4] in that no model is used to generate the turbulence; instead, it is generated by solving the filtered Navier-Stokes equations on a fairly large, laterally periodic domain, and letting the boundary layer and its turbulence naturally develop. Once the atmospheric boundary layer turbulence is developed, an array of actuator disks (as many as 108) is introduced into the flow field and the simulation continues with periodic boundaries in the horizontal direction. This approach, in effect, simulates an infinitely large wind plant. They perform fourteen different simulations in which they vary domain size, number of turbines, turbine thrust coefficient, and surface roughness. The aim of their work is to understand how a large wind plant modifies the vertical transport of momentum and kinetic energy across the atmospheric boundary layer, and to suggest a modified surface roughness that could account for the effect of a wind plant on the atmospheric boundary layer in regional- or global-scale weather simulations. Calaf et al.[6] do not make any comparison with field data because they are simulating an idealized wind plant and are more interested in the atmospheric physics than is Ivanell.[4]

Others have performed LES of smaller groups of turbines. Like Calaf et al.,[6] Porté-Agel et al.[7] created neutrally stable atmospheric inflow directly from LES. In this work, they compared different wind turbine models (actuator disk versus actuator line) and simulated a five turbine section of an operational plant and compare wake profiles with sonic detection and ranging (SODAR) data and power production predictions with supervisory control and data

acquisition (SCADA) output. Troldborg,[8] Mikkelsen et al.,[9] and Troldborg et al.[10] study the interaction of wakes created by an actuator line turbine representation, and they model the inflow in the same way as Ivanell.[4]

To provide a sound methodology for performing wind plant LES, this work combines the best aspects of the above-mentioned studies. Specifically, we wish to perform "engineering-focused" plant simulations, such as those of Ivanell[4] and Porté-Agel et al.[7] in which comparison to data from operational wind plants is made. We do not use a model for inflow turbulence, rather we generate atmospheric turbulence directly with LES, in the manner of Calaf et al.[6] and Porté-Agel et al.[7] We use an actuator line turbine model similar to that used by Porté-Agel et al.[7] and Troldborg,[8] which is described in Section II.B. Like Ivanell,[4] we locally refine the computational grid in the region of the turbines and their wakes. Unlike many of these researchers, the wind plant domain is not periodic, and the entire plant is modeled.

II. Physics Models

Two distinct physics models are used in performing our wind plant simulations: LES and the actuator line turbine aerodynamics model. LES solves the filtered Navier-Stokes equations in which body forces from the turbine model are imposed. The turbine model uses the velocity field from the LES to compute the aerodynamic forces imparted on the turbine blades.

A. Large-Eddy Simulation

With LES, the incompressible Navier-Stokes equations are spatially filtered to arrive at the resolved-scale (large-eddy scale) equations of motion. The formulation we use is similar to that of Moeng.[11] The filtered continuity equation is

$$\frac{\partial \widetilde{u}_j}{\partial x_j} = 0 , \tag{1}$$

where the overbar denotes filtering and $\widetilde{u}_j = u_j - u_j'$ is the resolved-scale velocity vector, which is the instantaneous velocity vector, u_j, minus the subfilter scale (SFS) velocity vector, u_j'. (Please note that a second-order accurate finite-volume formulation is used and a filter is not explicitly applied as is often done with pseudo-spectral calculations. Rather, the finite-volume formulation itself acts as a filter.)

The filtered momentum transport equation is

$$\frac{\partial \widetilde{u}_i}{\partial t} + \frac{\partial}{\partial x_j}\left(\widetilde{u}_j \widetilde{u}_i\right) = \underbrace{-2\varepsilon_{i3k}\Omega_3 \overline{u}_k}_{I} \underbrace{- \frac{\partial \hat{p}}{\partial x_i}}_{II} \underbrace{- \frac{1}{\rho_0}\frac{\partial}{\partial x_i} p_0(x,y)}_{III} \underbrace{- \frac{\partial \tau_{ij}^D}{\partial x_j}}_{IV} + \underbrace{g\left(\frac{\widetilde{\theta} - \theta_0}{\theta_0}\right)\delta_{i3}}_{V} + \underbrace{\frac{1}{\rho_0} F_i^T}_{VI} . \tag{2}$$

Term I is the Coriolis force due to planetary rotation in which ε_{ijk} is the alternating tensor, Ω_j is the rotation rate vector defined as $\mathbf{\Omega} = \omega[0, \cos(\phi), \sin(\phi)]$, and ω is the planetary rotation rate. Term II is the gradient of a modified pressure variable, $\hat{p} = \widetilde{p}' / \rho_0 + \tau_{kk}/3$, where $\widetilde{p}'(x,y,z,t) = \widetilde{p}(x,y,z,t) - p_0(x,y) + \rho_0 g z$. The quantity $\widetilde{p}(x,y,z,t)$ is the resolved-scale static pressure, $p_0(x,y)$ is the average static pressure at the surface that varies linearly in the horizontal directions due to a background driving pressure gradient, and $\rho_0 g z$ represents the hydrostatic variation of static pressure with height in which ρ_0 is the constant density of this incompressible flow. The gravitational constant, g, is taken to be 9.81 m/s^2 and z is height above the surface. The quantity $\widetilde{p}'(x,y,z,t)$, therefore, represents the deviation in resolved-scale static pressure from its time-averaged value. The modified pressure variable, $\hat{p}$, is density-normalized and lumped with the trace of the stress tensor, τ_{kk}. Term III represents the background driving pressure gradient. Term IV is the divergence of the deviatoric part of the fluid stress tensor, $\tau_{ij}^D = \tau_{ij} - \tau_{kk}\delta_{ij}/3$ (remember, the trace of the stress tensor is lumped with the pressure deviation scalar), where δ_{ij} is the Kronecker delta. The stress is composed of a viscous and a SFS part. Because this is a high Reynolds number flow, the SFS part dominates everywhere except near the surface. In the solver, the viscous part of the stress tensor is neglected everywhere except at the lower surface where the total (SFS plus viscous) stress is modeled using Moeng's surface stress model.[11] Term V is the Boussinesq approximation for buoyancy effects due to temperature variations in the flow in which $\widetilde{\theta}$ is the resolved-scale potential temperature and θ_0 is a reference potential temperature taken to be 300 K. In term VI, F_i^T is the body force field exerted by the actuator line turbine model.

The filtered potential temperature transport equation is

$$\frac{\partial \widetilde{\theta}}{\partial t} + \frac{\partial}{\partial x_j}\left(u_j \widetilde{\theta}\right) = -\frac{\partial q_j}{\partial x_j}, \tag{3}$$

where q_j represents the transport of temperature by molecular and SFS turbulence effects. As with the momentum transport equation's (Eq. 2) term IV, we neglect the molecular effects, except at the lower surface, because this is a high Reynolds number flow.

At all locations in the flow, except at the lower surface, the SFS stress and temperature transport are modeled with the linear gradient-diffusion relations

$$\tau_{ij}^{D} = -\nu^{SFS}\left(\frac{\partial \widetilde{u}_i}{\partial x_j} + \frac{\partial \widetilde{u}_j}{\partial x_i}\right), \tag{4}$$

$$q_j = \frac{\nu^{SFS}}{\mathrm{Pr}_t}\frac{\partial \widetilde{\theta}}{\partial x_j}. \tag{5}$$

The SFS viscosity, ν^{SFS}, is obtained with the standard Smagorinsky[12] model. The model constant, C_s, is set to 0.135. It should be noted that there are more sophisticated SFS models available and in use by the atmospheric community, such as the scale-dependent Lagrangian model.[13] The simplicity of the standard Smagorinsky model makes it attractive from an implementation point of view, but we plan to use one of these more sophisticated models in the near future. The turbulent Prandtl number, Pr_t, is generally 1/3, but is sensitized following Moeng's[11] work such that its value is closer to one in regions of local flow stability (i.e. where the vertical gradient of potential temperature is positive), such as the capping inversion at the top of the atmospheric boundary layer. In effect, the SFS temperature diffusivity is reduced in regions of stability.

We do not explicitly model sea surface waves, nor the fact that the sea surface has a velocity. Rather we model the lower boundary as a flat rough surface with an aerodynamic roughness height appropriate for the sea surface. At the lower boundary, τ_{ij}^{D} and q_j are directly specified using Moeng's[11] rough surface model. The inputs to Moeng's model are the friction velocity, u_*, and the vertical potential temperature flux at the surface. Monin-Obukhov similarity theory[14,15] is used to estimate u_*. At the upper boundary, τ_{ij}^{D} and q_j are set to zero.

B. Actuator Line Model of Turbine Blade Aerodynamics

We do not resolve the geometry of the turbine because the resolution required to perform true LES down to the geometry surface would be computationally infeasible. Therefore, we use Sørensen and Shen's[16] actuator line model of the turbine blades. The nacelle and tower of the turbine are not modeled.

Each actuator line rotates about the turbine's low-speed shaft, and the rotation rate is controlled using a combination of torque and pitch control as discussed in Section IV.B. Hence, the rotor a dynamical system that responds to changes in incoming wind. Each actuator line is divided into 40 equally spaced segments. The blade airfoil type, chord, twist, and local flow velocity are known at the center point of each segment. With that information and the rotor speed, the velocity magnitude, V_{mag}, and local flow angle, α_{local}, can be computed. The local flow angle is different from the angle-of-attack, α. The angle-of-attack is the angle between the airfoil chord line and the freestream, whereas the local flow angle is the angle between the airfoil chord line and the local flow. When an airfoil creates lift, there is an upwash in front of the airfoil and a downwash behind it due to the bound circulation. We assume that since the actuator line lies in the center of this bound vortex, the effects of up- and downwash are small, and we assume that $\alpha = \alpha_{local}$. Employing lift and drag look-up tables for the rotor airfoils, the magnitude of lift, L, and drag, D, at each actuator segment can be computed using

$$L = \frac{1}{2} C_l(\alpha)\, \rho_0 V_{mag}^2 c w \tag{6}$$

and

$$D = \frac{1}{2} C_d(\alpha)\, \rho_0 V_{mag}^2 c w\,, \tag{7}$$

where C_l and C_d are the lift and drag coefficients, respectively, c is the chord length, and w is the actuator segment width. Knowing that lift and drag forces act perpendicular and parallel to the oncoming wind vector, respectively, one uses those forces to form the total force vector experienced by the actuator segment, f_i^T.

Each actuator segment has its own value of f_i^T, which is a point force that cannot be directly applied to the flow field; rather, it must be smoothly projected onto the flow field volume. The volume integral of the projected force must be the same as the original actuator segment force. Following Sorensen and Shen,[16] the actuator forces are projected onto the computational domain as a body force field. At a location (x, y, z) in the domain, the body force field is related to the actuator segment forces using a Gaussian projection as follows:

$$F_i^T(x,y,z,t) = -\sum_{j=1}^{N} f_i^T(x_j, y_j, z_j, t) \frac{1}{\varepsilon^3 \pi^{3/2}} \exp\left[-\left(\frac{\left|\vec{r}_j\right|}{\varepsilon} \right)^2 \right]. \tag{8}$$

The summation is calculated over all N actuator segments of the turbine. The location of the j^{th} segment is (x_j, y_j, z_j), where $\left|\vec{r}_j\right|$ is the magnitude of the vector between (x, y, z) and (x_j, y_j, z_j), ε controls the width of Gaussian, and the negative sign accounts for the fact that the force that the turbine exerts on the flow field is equal and opposite to the force it experiences due to the flow. To maintain numerical stability, we set the quantity ε to twice the cube root of the volume of the cells local to the turbine. (The cells local to the turbine all have the same volume, so ε is a constant.) As pointed out by Troldborg,[8] this value of ε is roughly the minimum at which the force is smoothed enough to avoid spurious oscillations in the resulting velocity field using a central spatial discretization scheme, like we use.

The length ε should also correspond in some way to the actual airfoil chord length over which the lift and drag forces are computed. Therefore, it may make more physical sense to choose a value of ε such that the resultant Gaussian width is similar to the chord width. Since the Gaussian decays to 1% of its maximum value when $\left|\vec{r}_j\right| = 2.15\varepsilon$, choosing a value of $\varepsilon = c/4.3$ gives a Gaussian that roughly spans from the leading to the trailing edges of the airfoil. Because it would require a small grid spacing of $c/8.6$ near the actuator line to maintain stability with a central spatial-discretization scheme, though, we kept ε equal to twice the local grid size. Martínez et al.[17] have performed a fairly thorough analysis of the effects of the projection width.

III. Numerical Method

Our solver was created using version 2.0.x of the OpenFOAM (Open Field Operations and Manipulations) CFD toolbox.[18] The OpenFOAM software is a set of C++ libraries meant for solving partial differential equations. The governing equations are solved using the finite-volume method on unstructured meshes. All variables, except SFS quantities, are cell-centered and collocated on the grid. SFS quantities are located on cell faces because the divergences of SFS stress and temperature transport are ultimately required, and the divergence operator requires cell face quantities. This is in contrast to the standard OpenFOAM practice in which SFS quantities are solved at cell centers and then interpolated to cell faces to perform the divergence operation. With the OpenFOAM-standard method, we found the SFS friction near the surface to be excessively large, but direct computation of SFS quantities at cell faces remedies the problem. To avoid the pressure-velocity decoupling that occurs with collocated, incompressible solvers, the velocity fluxes at the finite-volume faces are constructed using an interpolation similar to that of Rhie and Chow.[19] All other interpolation from cell centers to faces is linear, which is similar to the second-order central finite-difference scheme. Time advancement uses Issa's[20] PISO (Pressure-Implicit Splitting Operation) algorithm, which is an implicit predictor/corrector scheme. The SFS quantities, turbine forces, and velocity fluxes within the convective term of the momentum equation are treated explicitly to avoid iterating during each time step, though. The implicit terms are integrated in time using second-order backward discretization. We use one predictor followed by three correctors. The momentum transport equation (Eq. 2) is solved directly. However, to enforce the continuity equation (Eq. 1), the divergence of the discrete momentum transport equation is taken, which results in an elliptic equation for the modified pressure. The momentum and potential temperature transport equations are solved using an iterative diagonal incomplete-LU pre-conditioned biconjugate-gradient linear system solver. The pressure equation, which is the most expensive to solve, is solved using a geometric

agglomerated algebraic multigrid solver. The code is parallelized using the message-passing interface (MPI), a necessity in problems of this size. Even producing the locally refined mesh required the use of parallel meshing tools. The domain is prepared for parallel processing using Scotch decomposition.[21]

IV. Simulation Methodology

The simulations are performed in two distinct phases: the atmospheric boundary layer precursor simulation and the wind plant simulation. In the precursor simulation, the atmospheric boundary layer and its turbulence are generated by the LES solver in a laterally periodic domain. In the wind plant simulation, the same LES solver is used, but the turbine aerodynamics model is activated, and the lateral domain boundaries are no longer periodic. Rather turbulence from the precursor simulation is fed into the upwind boundaries, whereas the downwind ones become outflow boundaries. Because each wind plant simulation is expensive, and a significant portion of this study is aimed at developing a sound wind plant simulation methodology, we only simulated one wind condition. Nonetheless, it was an expensive simulation.

A. Atmospheric Boundary Layer Precursor Simulation

The goal of the atmospheric boundary layer precursor simulation was to replicate conditions as closely as possible at Lillgrund. The atmospheric conditions there are outlined by Bergström.[22] We chose a mean hub-height (65 m) wind direction of 221.6° from the southwest and aligned with one of the turbine row directions. We also chose a mean hub-height wind speed of 9 m/s because it meant that the turbines would operate in Region 2, in which maximum energy is extracted from the wind creating the strongest wake deficits. According to Bergström,[22] this wind speed is typical of the 221.6° direction. The condition is of practical significance because wake effects will have a maximum effect on reducing the power production of the wind plant. Horizontally averaged wind speed and direction at a given height can be an input to the LES solver, and the driving pressure gradient vector (term III of Eq. 2) is adjusted each time step to maintain these conditions. The rotation rate vector was set to match that experienced at Lillgrund which lies at 55.52° N latitude.

More difficult to reproduce is a desired turbulence intensity

$$I(z) = \frac{\sigma_s(z)}{U(z)} = \frac{\sqrt{\overline{\langle u_s' u_s' \rangle}(z)}}{\sqrt{\langle \tilde{u} \rangle^2(z) + \langle \tilde{v} \rangle^2(z)}}, \tag{9}$$

where brackets denote horizontal averaging and an overbar denotes time averaging. The quantity $\sigma_s(z)$ is the velocity variance in the mean hub-height wind direction (the subscript "s" is meant to indicates a quantity in the "streamwise" direction), and $U(z)$ is the average resolved horizontal wind speed. A hub height turbulence intensity of about 6% is typical of the chosen wind direction.[22] The only other inputs to the LES solver (in addition to desired mean wind speed and rotation rate vectors) are surface aerodynamic roughness height, surface temperature flux, and the initial potential temperature profile, all of which affect the resultant turbulence intensity once the solution reaches a quasi-equilibrium state. Based on our past experiences, a fairly low turbulence intensity, like that desired, is obtained under neutrally stable conditions and low roughness. Therefore, we set the surface temperature flux to zero. A low surface aerodynamic roughness height of 0.1 mm was chosen by simultaneously solving Charnock's formula[23] relating surface stress to roughness height over water

$$z_0 = \frac{\alpha_c u_*^2}{g}, \tag{10}$$

in which the constant α_c is 0.016 as recommended by Garratt,[24] and the Monin-Obukhov scaling law[14,15] for neutral flows

$$\frac{U(z)}{u_*} = \frac{1}{\kappa} \ln\left(\frac{z}{z_0}\right), \tag{11}$$

where κ is the von Kármán constant set to 0.4. The initial capping inversion was created by specifying the initial potential temperature field to be 300 K up to 700 m above the surface. From 700 m to 800 m, the potential

temperature increases by 8 K. Above 700 m, the potential temperature increases at a rate of 0.003 K/m. This profile is similar to that used by Moeng and Sullivan[25] in simulating a neutrally stratified atmospheric boundary layer.

The domain size is 4 km in both horizontal directions and 1 km in the vertical direction. The resolution for the precursor simulation is a uniform 7 m in all directions requiring a mesh of 572 × 572 × 144 hexahedral cells. The lateral boundaries are all periodic. At the lower boundary, no condition on horizontal velocity is necessary because the surface stress is directly specified. The no-slip condition is not appropriate when a rough-wall-surface stress model is used. Simply, the velocity normal to the wall is set to zero. The same is true of temperature because a model for temperature flux is used on the lower boundary. On the upper boundary, the velocity normal to the boundary and the gradient of horizontal velocity normal to the boundary are zero. The temperature gradient there is set to match the initial temperature profile. On the upper and lower boundaries, the gradient of modified pressure is derived by taking the dot product of the momentum transport equation (Eq. 2) with the boundary normal direction. The initial velocity field was set uniformly to the desired hub-height wind vector with some small divergence-free perturbations near the surface.

The boundary layer was simulated for 12,000 s before it was considered to have reached a quasi-equilibrium state in which the initial transients had passed and the mean flow was undergoing gentle inertial oscillations. The simulation was run for an additional 2,000 s during which planes of velocity and temperature data on the south and west boundaries were saved each time step.

The domain was decomposed into 512 partitions for parallel processing. A variable time step in which the maximum Courant number of 0.75 was used. This corresponds to a physical time step of about 0.25 s, requiring roughly 37,000 processor-hours to compute the full 14,000 s of simulation time.

B. Wind Plant Simulation

The wind plant simulated in this study is the Lillgrund offshore facility[1] operated by Vattenfall Vindkraft AB. Lillgrund is located in the Öresund, a body of water between Copenhagen, Denmark and Malmö, Sweden. It is 7 km off the coast of Sweden and 7 km south of the Öresund bridge where the water depth is only 4 to 8 m. It consists of 48 Siemens SWT-2.3-93 three-bladed, upwind, horizontal-axis turbines[2] each with a rated power production of 2.3 MW, a rotor diameter of 93 m, and a rotor hub height of 65 m.

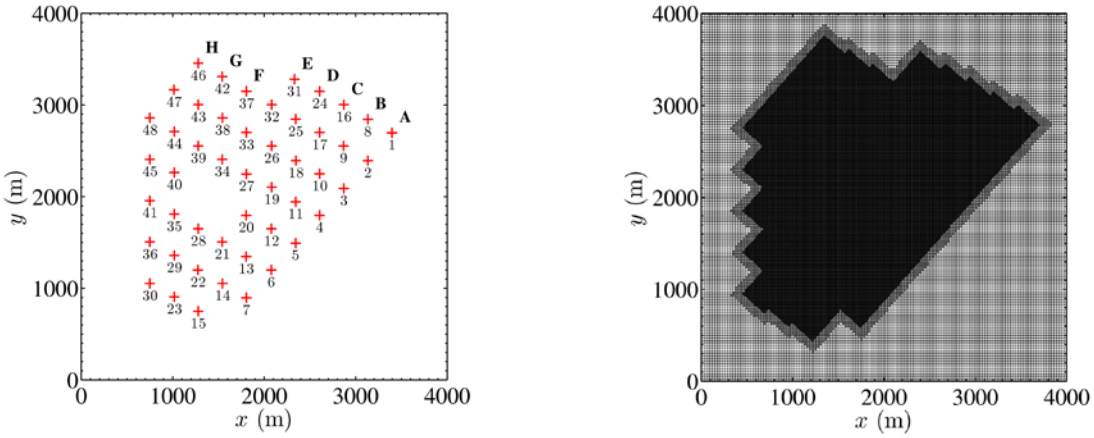

Figure 1. (a) A plan view of the Lillgrund wind plant simulation domain and the layout of the turbines in the plant. Positive *x* and *y* correspond to east and north, respectively. (b) A plan view of a horizontal plane of grid cells at turbine hub height showing the 7 m resolution out region, the 1.75 m resolution turbine/wake region, and the 3.5 m resolution buffer region.

Figure 1 is a plan view of the layout of the turbines at Lillgrund and with respect to the computational domain. Their positions were obtained from the report by Jeppsson et al.[1] Positive *x* and *y* corresponds to east and north, respectively. The turbines are placed 750 m from both the south and the west domain boundary. The mean hub height wind is aligned with southwest- to northeast-oriented rows, such as the one containing turbines 1 through 7.

Because the Siemens SWT-2.3-93 blade data are proprietary, we designed our own 2.3-MW turbine model to closely replicate the Siemens turbine. Based on information from the SWT-2.3-93 product brochure,[2] Porté-Agel et

al.,[7] Laursen et al.,[26] and Leloudas,[27] we were able to design a turbine rotor and analyze it with blade element momentum theory such that it has a power curve similar to that given in the SWT-2.3-93 product brochure. Figure 2 shows that we achieved reasonable agreement with the manufacturer's published power curve, especially in the region of wind speeds around 9 m/s. According the SWT-2.3-93 product brochure,[2] NACA63 series and FFA airfoils are used in the actual blade. We used members of those airfoil families in our turbine model, and the two-dimensional lift and drag data were obtained from Bertagnolio et al.[28] Using the National Renewable Energy Laboratory's (NREL) AirfoilPrep code,[29] the two-dimensional lift and drag data corresponding to stalled conditions were corrected for the fact that stall may be delayed on rotating blades using the methods of Du and Selig[30] and Eggers et al.[31] AirfoilPrep was also used to apply Viterna and Janetzke's method[32] to extrapolate the data to a wider range of angles of attack. We also designed a simple torque controller similar to that of the NREL 5MW reference turbine model[33] that contains different control strategies for Regions 1, 1-1/2, 2, and 2-1/2. To provide the smooth transition of the power curve into Region 3, which is shown in the manufacturer's published power curve, the blades begin to pitch below the rated wind speed of 13–14 m/s. With this torque and blade pitch control, the rotor speed is governed by a dynamical system that reacts to the oncoming wind, the aerodynamic torque the rotor creates, and the opposing torque created by the generator.

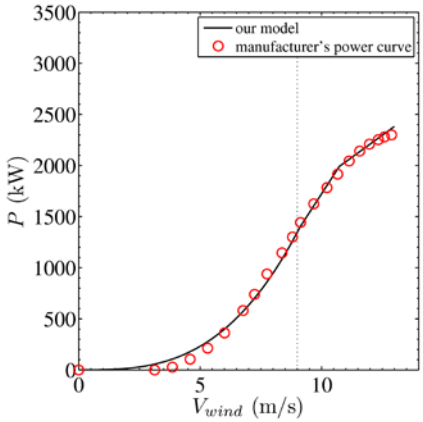

Figure 2. The power curve of our 2.3-MW turbine model (black line) as compared to manufacturer's data[2] for the Siemens SWT-2.3-93 wind turbine (red circles).

No longer are periodic boundary conditions used, as were used in the precursor simulation. The upper and lower boundary conditions remained the same, but the velocity and temperature at the south and west boundaries are specified by using saved planes of data from the precursor simulation from those boundaries. The gradients of velocity and potential temperature normal to the north and east boundaries are set to zero with the condition that flow must only exit the domain. The velocity predicted on these boundaries using this method is then adjusted such that the total flux of velocity on all domain boundaries is zero. The gradient of pressure normal to all boundaries is set by using the dot product of the momentum transport equation (Eq. 2) and the boundary normal direction. The driving pressure gradient was set to a constant value that is the average of the time-varying driving pressure gradient used in the precursor simulation. The average was only taken over an interval in which the flow had reached a quasi-equilibrium state.

In creating the computational mesh for the wind plant simulation, the 7 m resolution mesh from the precursor simulation was used as the "background" mesh. The mesh was then locally refined around the turbines and in their wakes by splitting the hexahedral background mesh cells in each direction until 1.75 m resolution was obtained. Because of this local refinement, the wind plant mesh contains about 315 million cells.

The domain was decomposed into 4096 partitions for parallel processing. The time step in the wind plant simulation is governed by the rotor tip speed. We desire to limit the movement of the rotor tip to a grid cell length each time step. This requires that the time step was fixed to 0.015 s. To simulate ten minutes of wind plant flow, roughly one million processor-hours were required.

V. Results

A. Atmospheric Boundary Layer Flow

Figure 3 shows the vertical profiles of horizontal- and time-averaged wind speed and turbulence intensity from the atmospheric boundary layer precursor simulation. As is typical of neutrally stratified flows, as opposed to an unstably stratified flow, there is significant vertical shear across the rotor disk. In this case, the wind speed increases by 1.38 m/s vertically across the disk. Turbulence intensity at the rotor hub height is 6.2%, which agrees well with the reported conditions at Lillgrund for this wind direction.[22] Turbulence intensity decreases with height above the surface.

The mean vertical profiles shown in Fig. 3 do not give information about the structure of the turbulence in the wind. Contours of instantaneous streamwise and vertical velocity fluctuations about the mean hub height value taken in a horizontal plane at the hub height are shown in Fig. 4. In the streamwise flow, there are significant elongated low- and high-speed structures aligned with the flow. These structures become important because turbine rotors can sometimes be situated inside theses structures and at other times be situated outside of them. When inside such structures, power production will decrease or increase. Other atmospheric LES modelers, such as Moeng and Sullivan[25] and Khanna and Brasseur[34] have also observed such elongated structures when simulating the neutral atmospheric boundary layer. The instantaneous vertical velocity fluctuations are much smaller and lack the structure of those in the horizontal. This is characteristic of neutral flow in which there are no buoyancy-induced vertical motions.

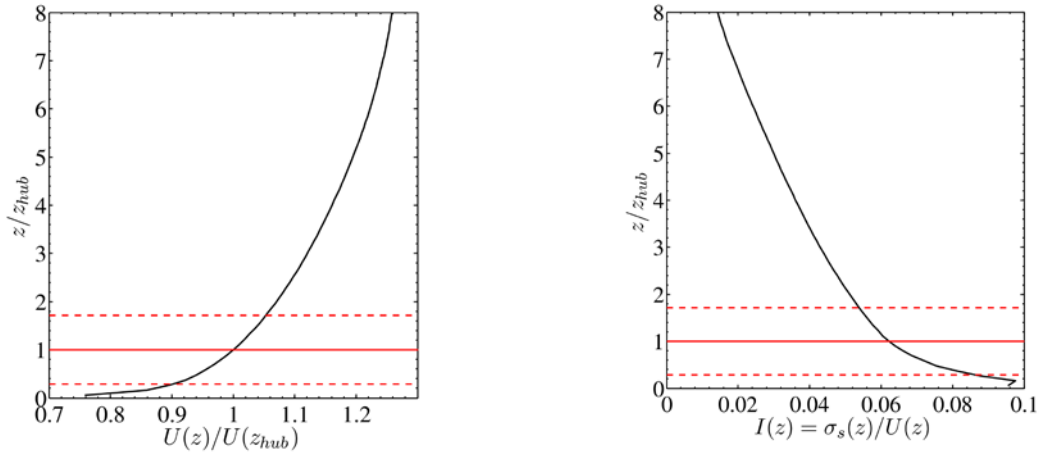

Figure 3. Vertical profiles of the horizontally and time-averaged (a) velocity and (b) turbulence intensity from the precursor simulation. The solid red horizontal line is at hub height. The dashed red lines indicate the vertical extent of the turbine rotor.

B. Wind Plant Flow

Figure 5 (a) shows contours of instantaneous resolved streamwise velocity normalized by mean hub-height wind speed taken on a horizontal plane at the turbine hub height from the wind plant simulation. The low-speed wakes are clearly visible behind the 48 turbines. In some places, the wake velocity magnitude is only about 50% of that of the freestream wind. Significant meandering of the wakes is visible, especially behind turbines of the sixth row or further downstream. We will show later that power production is slightly enhanced by this meandering as compared to when there is little or no meandering. With the increased meandering, turbines are more often placed in a partially waked situation than in the little- or no-meandering situation when turbines are almost always fully waked. Increased meandering, and increased partial waking, also creates more damaging fatigue loads as turbine blades periodically pass into and out of wakes. Also, within the refinement region, one can observe smaller turbulent scales than are possible to resolve on the outer grid. How much distance it takes for the cascade of resolved turbulence from the less refined to the more refined grid to occur is unclear, but should be studied in future work.

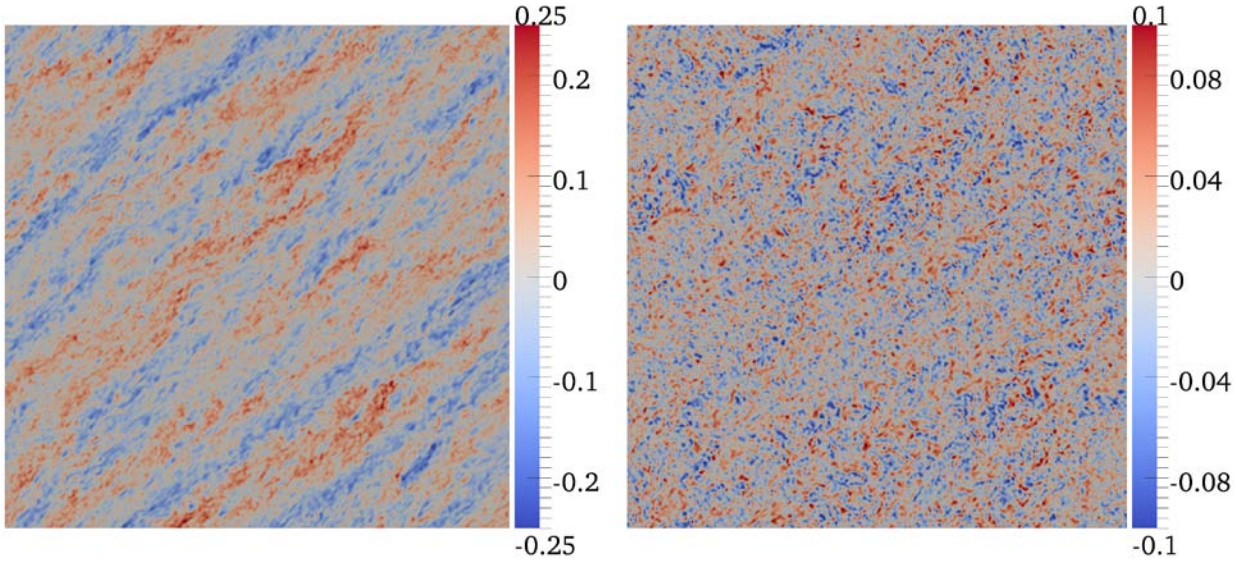

(a) (b)

Figure 4. Contours taken in a horizontal plane at the rotor hub height during the precursor simulation of instantaneous (a) horizontal and (b) vertical velocity fluctuations from the mean normalized by the hub height wind speed.

Figure 5 (b) shows contours of mean streamwise velocity normalized by mean hub-height wind speed in the same horizontal plane. The mean is taken over ten minutes after the flow becomes fully developed. The wakes are clearly visible. Near the center of the rotor, there is a higher speed region of flow since the blades are not so heavily loaded there, and we do not model the effect of the nacelle to decelerate the fluid. The region of higher-speed fluid persists for a longer distance behind the front turbines than all other members of a row. This is because the increased turbulence in the wakes of the downstream turbines more quickly mixes momentum across the span of the wakes. The wakes of the last turbines in each row persist at least to the downstream boundary. It would be useful, in terms of understanding the interaction of wind plants situated near other wind plants, to extend our downstream domain boundary and quantify the distance over which the wake of the wind plant is significant.

Figure 5 (c) depicts contours of resolved turbulent kinetic energy normalized by the hub-height wind speed in this same hub-height horizontal plane. Resolved turbulent kinetic energy is greatest near the edges of the wake where there is most shear to generate turbulence. Resolved turbulence kinetic energy appears greatest in the wakes of the further downstream turbines. The meandering of these further downstream wakes likely is the cause of this increased turbulent kinetic energy. This brings up the important point that the resolved turbulent kinetic energy calculated in a stationary reference frame is different than that calculated in a reference frame meandering laterally and vertically with the wake, but at a fixed streamwise location. Understanding the difference between turbulent kinetic energy calculated in these two different ways may be useful for those attempting to model wind plants using the Reynolds-averaged Navier-Stokes (RANS) equations. A RANS turbulence model designed for wind plant flow may include a source of turbulence created by the rotor blades, and a separate source due to meandering.

Downstream of the last row of turbines, discrete decreases in resolved turbulent kinetic energy can be seen at locations where the flow exits the local grid refinement regions. On a coarser grid, it is expected that there will be less resolved turbulent kinetic energy and more subfilter scale turbulent kinetic energy since fewer turbulent length scales are resolved there. There are also small patches of resolved turbulent kinetic energy at the upstream side of the local refinement region. It is possible that there is some localized oscillation in the flow field spuriously causing these patches of increased resolved turbulent kinetic energy. The oscillation could be due to the abrupt change in grid size, and hence filter width. In the future, we may try smoothly varying the filter width across the boundary of the locally refined region of the grid.

Figure 5 (d) shows contours of instantaneous coefficient of modified pressure (pressure with the effects of the driving pressure gradient and the vertical hydrostatic variation subtracted out), where the velocity used in the normalization is the mean hub-height wind speed at the turbine hub height. The most noticeable feature of the pressure contours is that pressure is higher upstream of a turbine than downstream of one, which is to be expected. Also interesting is that a slight negative pressure gradient has formed across the wind plant. This simulation was run

with a background driving pressure gradient set to the average value used during the atmospheric boundary layer precursor simulation. It appears that a slightly stronger gradient is required for the wind plant simulation to overcome the additional drag on the atmospheric boundary layer caused by the turbines. Also, there are significant pressure fluctuations near the edges of the wakes, indicative of turbulent motions and structures.

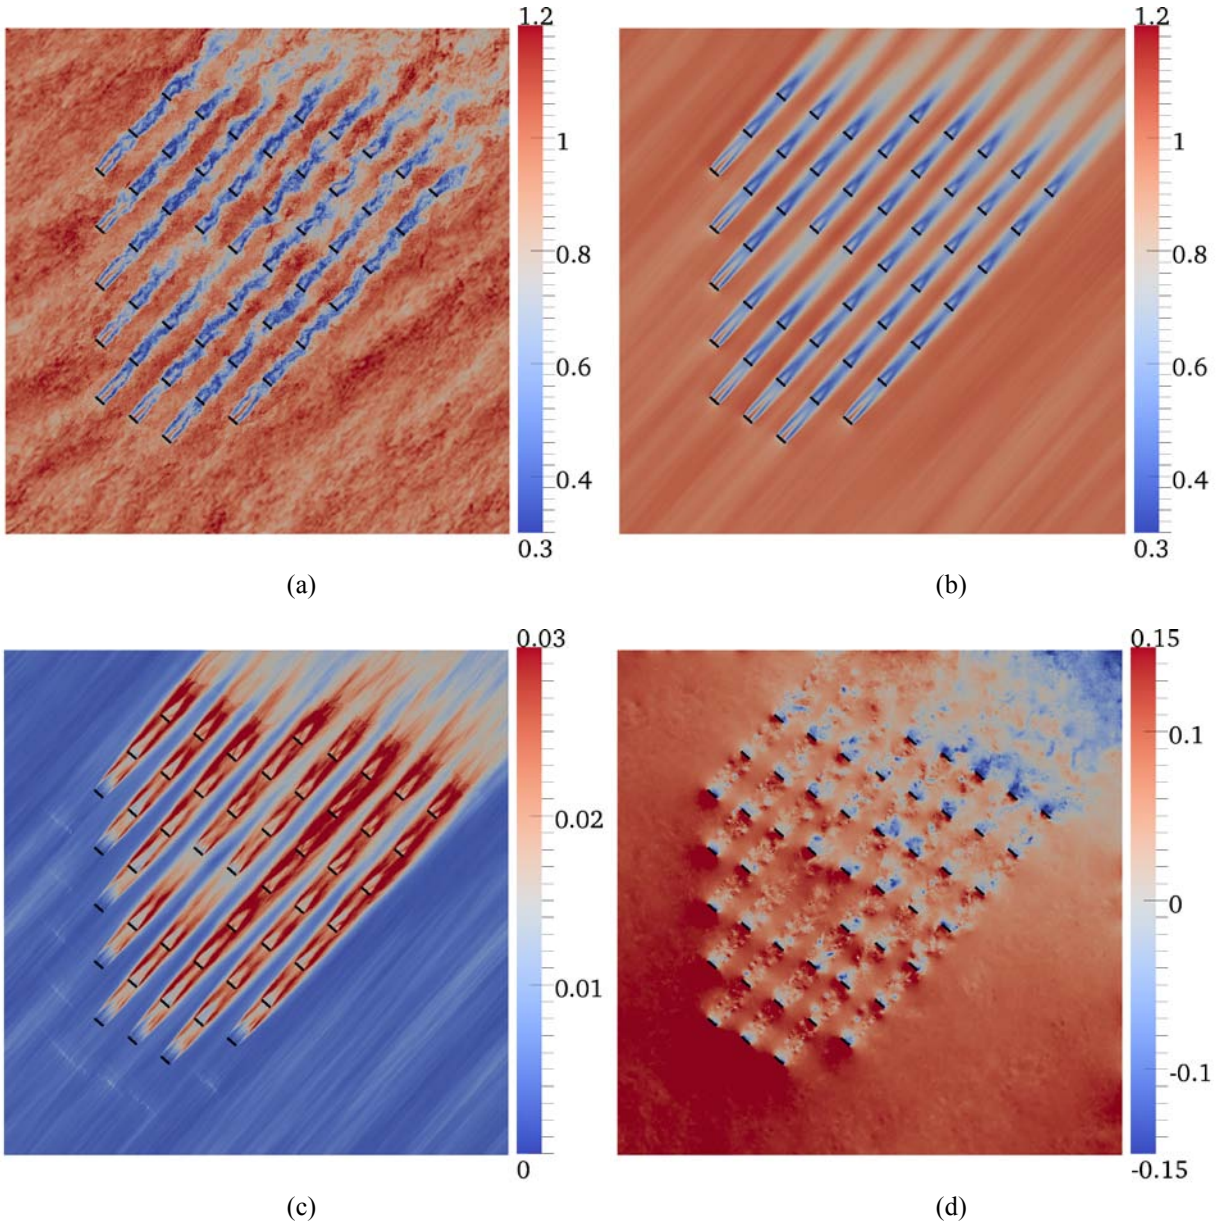

Figure 5. Contours taken in a horizontal plane at the rotor hub height of (a) instantaneous and (b) time-averaged resolved streamwise resolved velocity normalized by hub height wind speed, (c) resolved turbulent kinetic energy normalized by hub height wind speed squared, and (d) instantaneous resolved modified pressure normalized by $1/2\rho_0 U_{hub}^2$. The black bars indicate the turbine rotors.

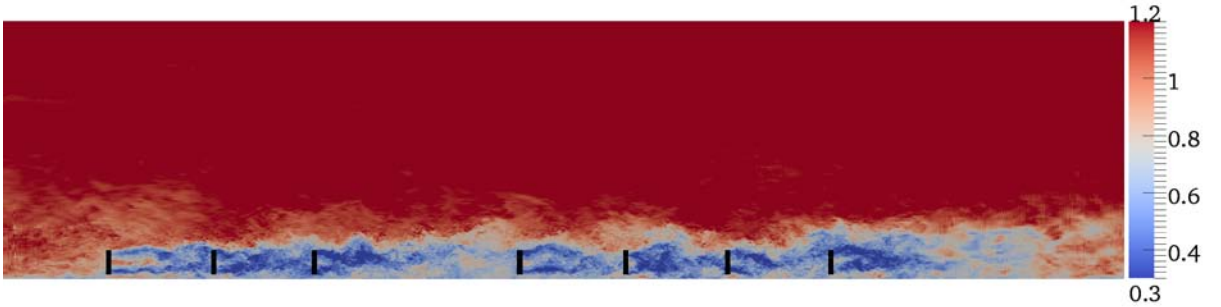

(a)

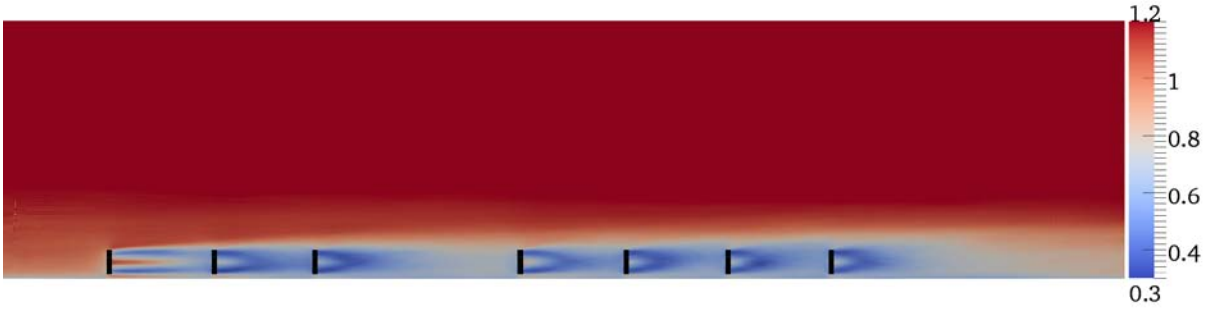

(b)

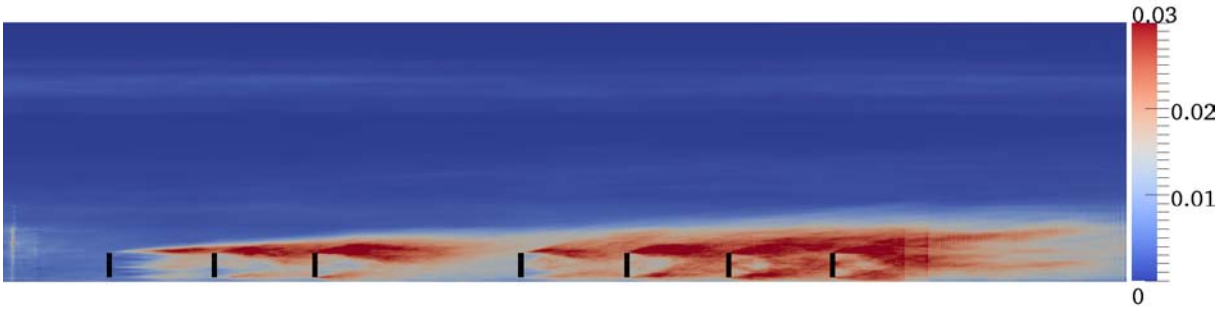

(c)

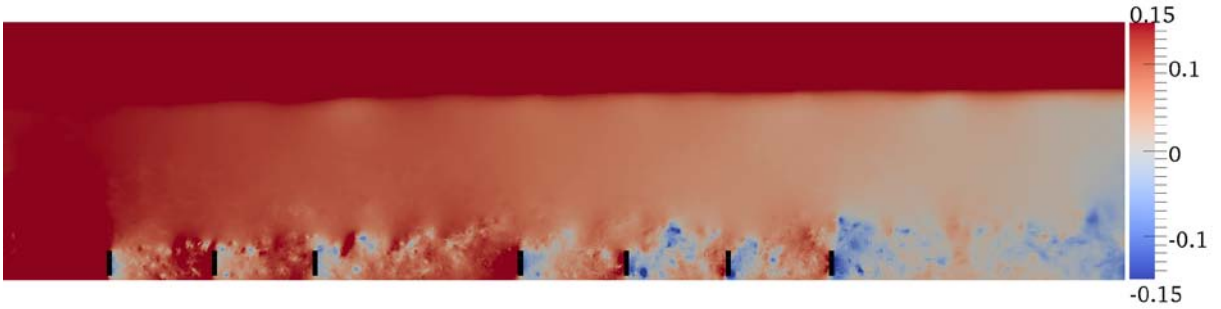

(d)

Figure 6. Contours taken in a vertical plane through row D of (a) instantaneous and (b) time-averaged resolved streamwise velocity normalized by hub height wind speed, (c) resolved turbulent kinetic energy normalized by hub height wind speed squared, and (d) instantaneous resolved modified pressure normalized by $1/2\rho_0 U_{hub}^2$. The black bars indicate the turbine rotors.

Figure 6 shows all the same variable as in Figure 5, but the contour plane is vertical and taken through the center of the turbines in row D (refer to Figure 1 (a) for row lettering). Most of the same behavior seen in the horizontal contour planes is seen here. Interestingly, Figures 6 (b) and (c) indicate the existence of an internal layer that forms over the wind plant. In Figure 6 (b), one can observe a gradual increase of height of the velocity deficit associated with the wakes with downstream distance. Figures 6 (c) shows a more noticeable increase in height of the region of increased turbulent kinetic energy with downstream distance. Chamorro and Porté-Agel[35] observed the formation of such layers in a wind tunnel study using miniature wind turbines. They liken the layer to that which forms downstream of a smooth to rough surface transition. Calaf et al.[6] also observed such layers in their simulations.

One of the main goals of this work is to explore the ability of this methodology to predict power production. Production data from Lillgrund are presented by Dahlberg,[3] and we compare to that data in Fig. 7. Figure 7 shows the time-averaged power produced by each turbine, $\overline{P}_i$, normalized by the time-averaged power of the corresponding first turbine of the row (rows B–D are normalized differently, as discussed in the next two paragraphs). The average is over a ten-minute interval that starts after the initial transients of wake development have passed, which required roughly 150 s. The power plots are separated by the different wind-aligned turbine rows. Turbine spacing along these rows is 4.3 rotor diameters, which is very close by today's standards. Each row is laterally separated from adjacent rows by 3.3 diameters. The effect of wakes on power production is dramatic. The second through fifth turbines in the row produce only 30–40% of the power of the first turbine in the row. After the fifth row, there appears to be some power recovery to 40–50% of the first turbine average power. This may be due to the increased wake meandering of the more downstream wakes observed in the instantaneous horizontal velocity contour shown in Fig. 5 (a). The increased meandering means that the turbines further downstream spend more time only partially waked as compared to the more upstream turbines that do not experience as high a degree of wake meandering and spend more time fully waked. There is a marked increase in power production results for turbines 27 and 34 in rows D and E, respectively, compared to the turbines surrounding them. Those turbines produce 50–60% of the time-averaged power of the first turbines in rows D and E. Inspection of Fig. 1 (a) shows that a turbine is missing immediately upstream of both turbines 27 and 34. When Lillgrund was built, those turbines were omitted because of inaccessibility to construction ships due to the shallow water depth there. The effective upstream spacing for turbines 27 and 34 is then 8.6 rotor diameters, the cause of the much improved performance.

Dahlberg[3] presents time-averaged power production results for turbines in rows B–D. These results are taken from a database of slightly more than a year's worth of data. Although the data were binned by wind direction, and we compared only with the down-the-row direction simulated in this study, the field data is a composite of many wind speeds (all below the turbine's rated power production wind speed) and turbulence levels. As reported by Bergström,[22] the averages of these wind speeds and turbulence levels are similar to the single condition we simulate. Therefore, a perfect comparison would only be possible if we were to run simulations of the many different wind conditions that occurred at Lillgrund during the data collection period and then average those results together. Such a task would be difficult with our computer resources. Nonetheless, a comparison can be made between our single ten-minute average and Dahlberg's year of data.

It is unclear how Dahlberg[3] normalized the time-averaged power production data for turbines within rows B–D. There appear to be two options: 1) he normalized by the time-averaged power production of the first turbine in the corresponding row (for example, he normalized time-averaged power produced by turbines in row B by that of turbine 15), or 2) he normalized by the mean of the time-averaged power production of the first turbines in rows B–D (turbines 15, 23, and 30). Since the time-averaged normalized power that he reports for the first turbines of these rows are close but not equal to one, it seems that Dahlberg normalizes using option 2). For rows B–D, we follow Dahlberg's normalization; however, for all other rows, we normalize by the time-averaged power of the first turbine in the corresponding row.

Superimposed onto the row B–D data in Figure 7 is the Lillgrund field data (shown in red) presented by Dahlberg.[3] Figure 8 shows the absolute difference between the predicted time-averaged normalized power and the Lillgrund data. Generally, for all three of these rows, the most error is observed with the sixth and higher turbine downstream. With those turbines, differences between the simulated and observed normalized time-averaged power production of positive 0.07 to 0.12 are observed. This means the LES overpredicts the time-averaged power for these turbines by 25–40%, which is significant. These downstream turbines encounter wakes in which we observe the greatest amount of meandering. Possibly, our LES model is overpredicting the amount of meandering present in the further downstream wakes. Also plausible is that our single wind condition is accurately allowing such meandering to occur, but the many other wind conditions that occurred over the year of field data collection did not allow the same meandering to occur, and we do not capture those other conditions. Also, we do not model turbine nacelle yaw control; possibly, this simplification is partially to blame for the overprediction.

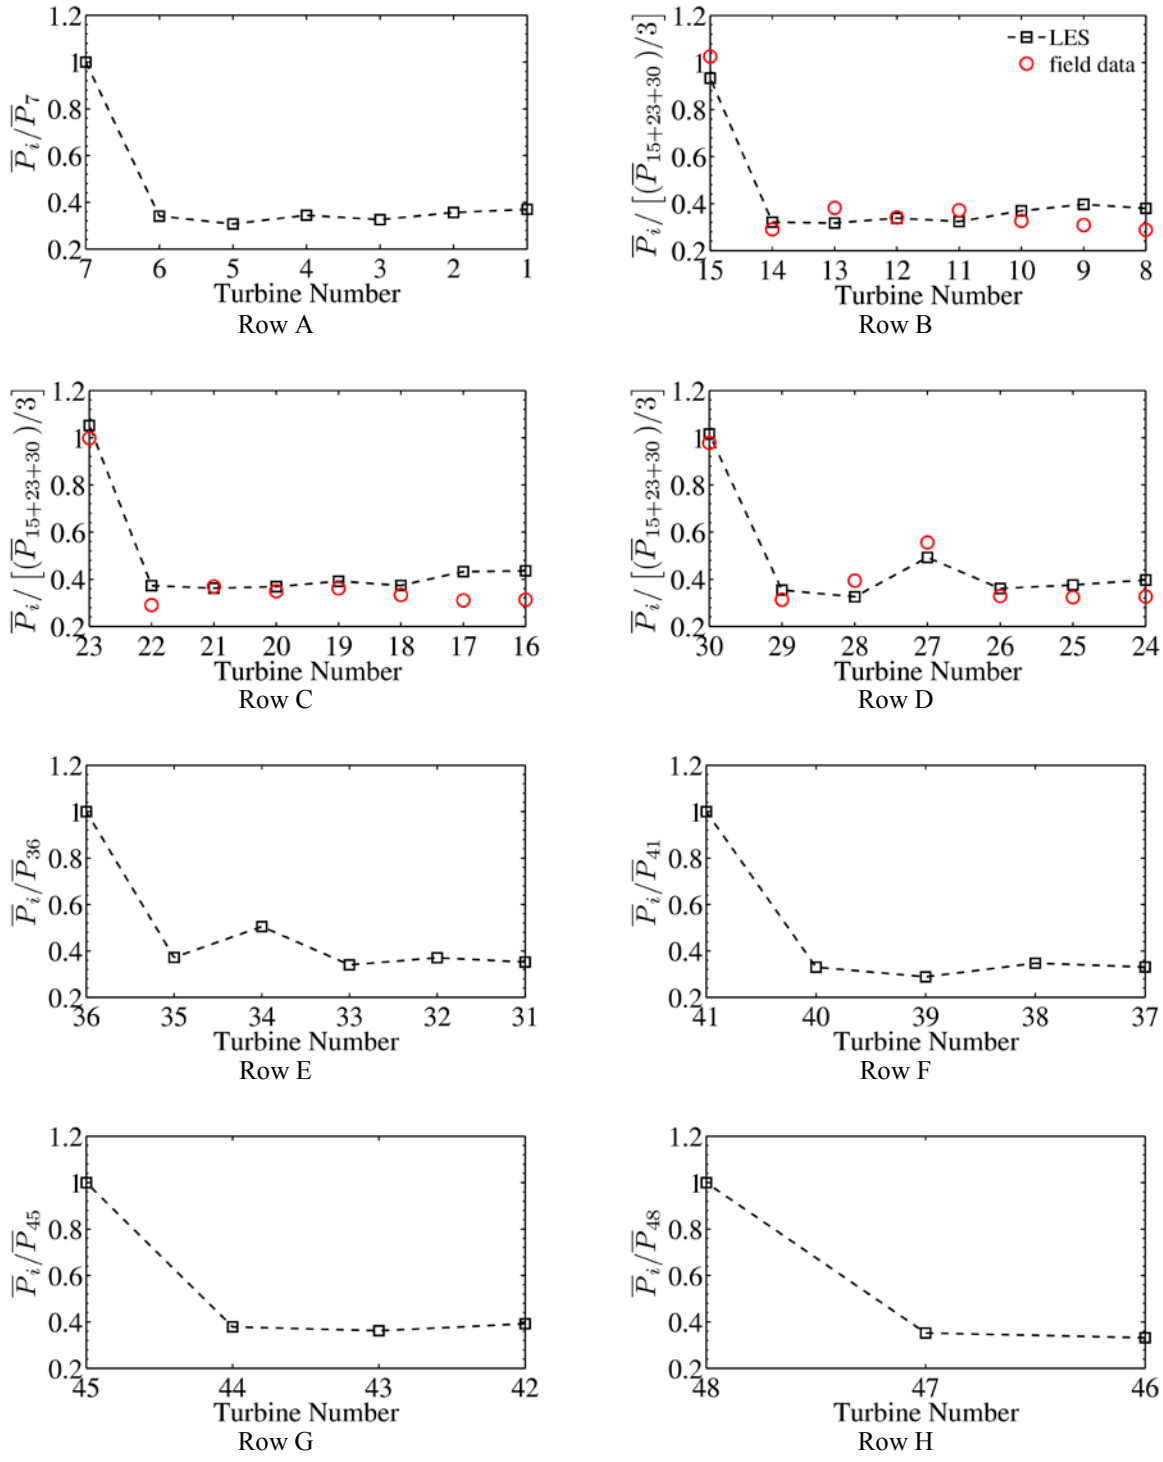

Figure 7. Simulated time-averaged power produced by each turbine, $\overline{P}_i$, normalized by the average power of the first turbine in the row (black squares with dashed line). Comparison with Lillgrund production data[3] (red circles) is shown for rows B–D.

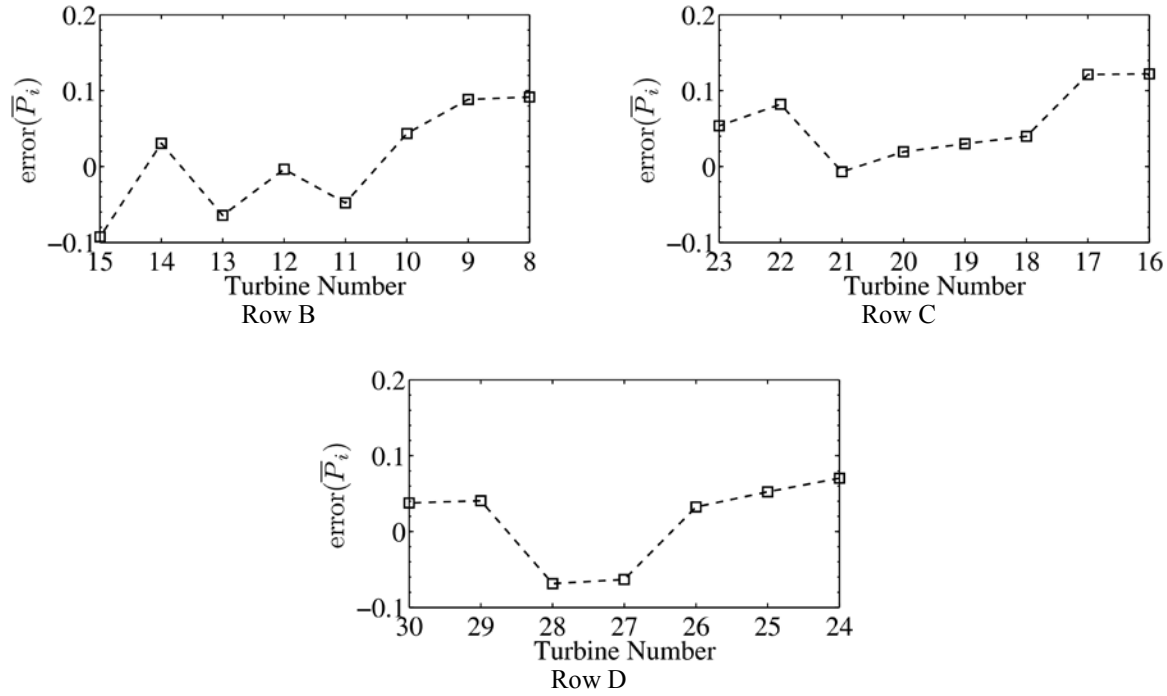

Figure 8. Difference in time-averaged power predicted by the LES and the Lillgrund production data[3] for rows B–D.

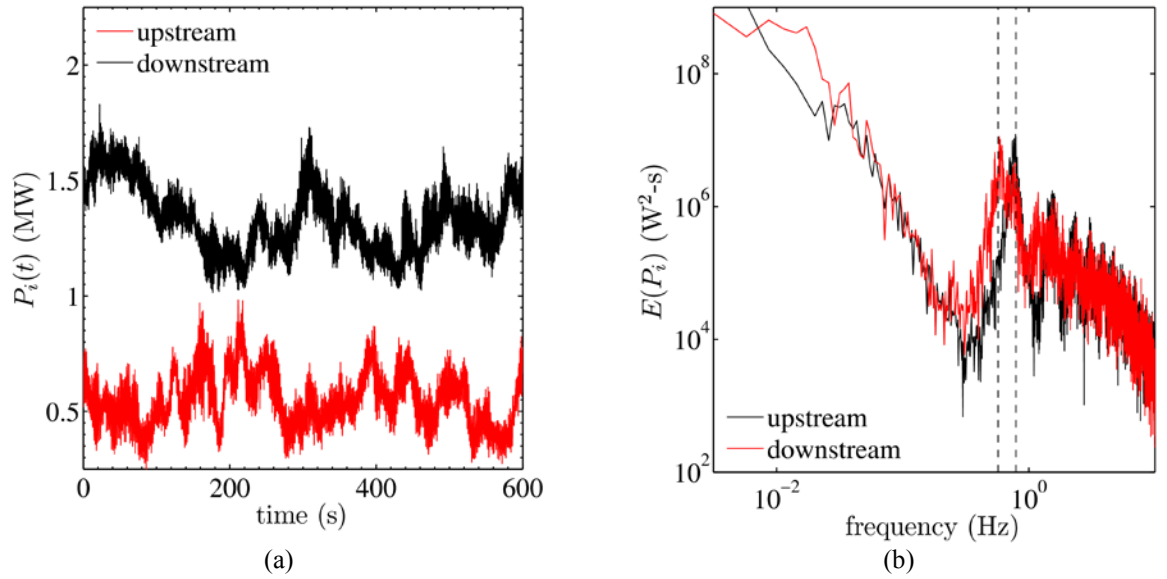

Figure 9. Plot (a) shows sample time histories of power produced by the turbine at the downstream (turbine 8 – red) and the upstream end (turbine 15 – black) of row B. Plot (b) shows a spectrum of these two time histories.

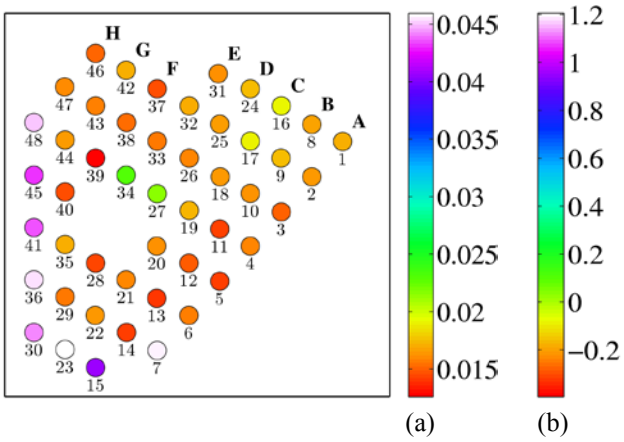

Figure 10. An overview of each turbine's time-averaged power production. Scale (a) shows the fractional contribution of each turbine to the entire plant's power production. Scale (b) shows the fractional deviation of each turbine's time-averaged power production from the time-averaged plant power production.

Since we have been examining time-averaged power production data, it is informative to study the time-varying data. In Fig. 9 (a), we show the time history of power produced by the most upstream and downstream turbines of row B, turbines 15 and 8, respectively. The most noticeable feature of these time histories, aside from the difference in power levels, is that there appear to be both low- and high-frequency variation in the signal. Fig. 9 (b), shows a frequency spectrum of the two power time histories. There is a clear peak (highlighted by a vertical dashed line), in each spectrum at a frequency corresponding to the blade-passage frequency, which is three times the rotor rotation frequency on these three-bladed turbines. For the upstream turbine 15, this frequency is 0.79 Hz, and for the downstream turbine 8, the frequency is 0.57 Hz. These correspond to rotor rotation rates of 15.8 and 11.4 RPM, respectively, which agree with the simulated rotation rates for these turbines. The upstream turbines' rotors rotate faster than the downstream ones because they are subject to higher wind speeds, and these turbines are variable speed machines. This brings up the point that in a simulation of this type, it is important to include variable-speed control, as we have done, to more accurately predict power. The downstream turbine's power production spectrum shows higher energy content in frequencies in the range of 0.008–0.026 Hz. This frequency range, which corresponds to time periods of 40–125 s, could be due to wake meandering. Inspection of Fig. 5 (a), which shows an instantaneous view of resolved velocity in a horizontal plane at hub height, shows that the length of one period of wake meandering is roughly the turbine spacing of 4.3 rotor diameters. If the wakes propagate at 0.4–0.6 times the mean flow speed, then the timescale of one wake meandering period is 75–110 s, corresponding well to the higher energy low-frequency range observed in the downstream turbine's power production spectrum.

Ultimately, in terms of sending energy to the electrical grid, the interest lies in how the turbines perform collectively as a power plant. For this wind direction, we computed the efficiency of the entire wind plant by dividing the sum of the time-averaged power of each turbine by the mean of the time-averaged power of the front turbines in rows B–D (turbines 15, 23, and 30). Again, Dahlberg[3] does not state which turbines he uses as "isolated" turbines to compute plant efficiency, but it appears that our method is consistent with his. Our simulation gives a plant efficiency of 47.6% for this wind direction, and Dahlberg reports an efficiency of 46.3%.

We can examine the fraction of total time-averaged power that each turbine produces. We can also study the deviation of each turbine's time-averaged power from the mean of the time-averaged power production of all the turbines in the plant. Figure 10 shows these quantities—the distribution of these quantities amongst the turbines is the same, but the magnitudes of the quantities are different. Color scale (a) shows the fractional contribution of each turbine to total power, and color scale (b) shows the deviation from the average. Each turbine at the front of a wind-aligned row produces 4–4.5% of the total power, whereas the downstream turbines produce 1.5–2% of the total power. Turbines 27 and 34, which lie behind the omitted turbines, produce about 2.2% of the total power. There is a clear variation in power produced by turbines at the same position in different rows, like the front row turbines. The variance of the front row turbines is 3.7% and that of all others (except turbines 27 and 34) is 8.3%. We expect

that if our averaging time were greater than ten minutes, this variation would diminish. The downstream turbines produce between 10 and 40% less power than the mean of time-averaged power produced by all turbines in the plant, whereas the front row turbines produce 100–120% more power.

VI. Conclusion and Future Work

We have presented aerodynamics and power-production results from a large-eddy simulation of the Lillgrund offshore wind plant, which contains 48 2.3-MW Siemens turbines. This simulation is an important step towards establishing a "best practices" methodology for performing wind plant large-eddy simulation, and it also identifies problems to be further researched. The scale of such a computation is large, requiring about one million processor-hours to obtain ten minutes of data. A simulation of this type is not meant to be a replacement for faster, computationally lighter engineering tools, but rather a research tool used to augment field observations in gaining a better understanding of wind plant aerodynamics and power production. The tool helps us understand wake creation, propagation, and interaction with other wakes and the atmospheric boundary layer. It also helps us understand the dependency of power production on aerodynamics. Furthermore, in other work by the authors,[36] we have made progress in coupling this tool with a wind turbine structural and system dynamics model so that we can explore structural response to atmospheric turbulence and wakes, and control strategies to mitigate damaging loads and improve power production. The hope is that knowledge gained from such simulations can be used to greatly improve computationally lighter wind turbine and wind plant analysis tools used by engineers and planners in predicting quantities such as annual energy production and damage equivalent loads.

Our simulation predicts time-averaged turbine power production well up to the fifth turbine in a wind-aligned row. Beyond that, the predicted wake meandering produces relative power overprediction errors of 25–40%. This is in contrast to Ivanell's[4] simulations of the Horns Rev wind plant, in which power underpredictions of about 30% were observed when wind is aligned with the turbine rows. We believe that our use of deterministic inflow turbulence created by an atmospheric large-eddy simulation on a domain kilometers in width creates large-scale atmospheric motions that partially drive wake meandering; stochastic turbulence generation methods, such as that used by Ivanell, possibly may not reproduce as strong a meandering effect. It is this meandering that we think increases power production of downstream turbines by partially waking them, as opposed to fully waking them, for a larger percentage of time. Our simulation, though, predicts the overall wind plant efficiency within one percent.

The difficulty in making comparisons with field data became evident in this study. Publically available field data, like that with which we compared, is often binned by wind direction and averaged over a very long time—in this case, over a year's time span. Our simulation was only ten minutes in duration. The average wind speed and turbulence intensity of the wind coming from the southwesterly wind direction we simulated has been measured at Lillgrund, and our simulation well matched those average conditions. However, the reality is that the field data is composed of many wind conditions from a single direction with a mean similar to our one condition. In addition to the neutrally stratified condition we simulated, the field data undoubtedly contains conditions, such as stable and unstable stratification, that produce different flow phenomena and behavior that become part of the average. In order to make a better comparison, we need to run more case with different atmospheric conditions.

This simulation highlights the fact that in a wind plant with spacing as small as at Lillgrund (4.3 rotor diameter streamwise spacing for the wind direction studied here), the turbines behind the first row incur a significant power production penalty. These turbines only produce 30–40% of the power of the first row turbines. It will be an interesting future study to examine the structural response of the turbines in the plant.

We also see that the wind plant creates an internal velocity deficit and turbulent kinetic energy layer similar to that observed experimentally by Chamorro and Porté-Agel[35] and computationally by Calaf et al.[6] The exchange of momentum of this layer with the rest of the atmospheric boundary layer above and to the sides of the wind plant would be an interesting study. The results of such a study could be used to design better RANS turbulence models for wind plants, or even reduced-order models to be used with mesoscale weather models. Behind the wind plant is a train of wakes from each turbine row that extend at least to the downstream domain boundary. We should extend this domain boundary to study just how far the wake of a wind plant persists. That knowledge becomes important when multiple wind plants are situated near one another.

There are many unresolved issues that we identified in performing this simulation. One of the larger issues is that of boundary condition types and locations. We are not sure how far away from the wind plant the side and top boundaries should be placed such that they do not spuriously influence the flow. We currently use a rigid lid upper boundary condition, and in reality, flow may desire to exit this boundary as it passes over the wind plant. A

convective boundary condition may be more suitable. Another boundary condition-related issue stems from Moeng's[11] surface stress model applied to the lower surface. It is meant for horizontally homogeneous flow and therefore relies upon horizontal averages. In wind plant flow, horizontal averages are no longer appropriate. Running time, local spatial, or Lagrangian averages along a surface streamline may be better options. This issue will also arise when terrain is introduced. Another important issue is that of subfilter scale modeling. We use the standard Smagorinsky model, and there are more sophisticated models available and in use by the atmospheric community, such as the scale-dependent Lagrangian model.[13] Implementation of such a model would allow us to pursue the more challenging stably stratified condition. In Section II.B, we discuss the need for improvement in the projection of turbine aerodynamic forces onto the flow field. Last, adaptive mesh refinement would be useful in providing higher resolution only where necessary, but may incur a run-time penalty in performing the refinement and the processor load balance.

Acknowledgments

The simulations performed in this work were computed on the Red Mesa high performance computing system of the National Renewable Energy Laboratory. W. Jones and M. Bidwell helped us to acquire the compute time to run these simulations and A. Purkayastha assisted in compiling OpenFOAM on Red Mesa. P. Spalart continues to provide extremely insightful thoughts about actuator line modeling. E. Paterson has always been available to help overcome OpenFOAM-related problems. M. Lawson provided useful thoughts on large-scale grid generation strategies. P. Sullivan and E. Patton both helped overcome problems with computing the stress term in our solver. M. Robinson has been instrumental in defining questions to be addressed in wind plant modeling. Funding for this work is from the U.S. Department of Energy.

References

[1]Jeppsson, J., Larsen, P. E., and Larsson, Å., "Technical Description Lillgrund Wind Power Plant," [online report] Vattenfall Vindkraft AB, 2_1 LG Pilot Report, Sept. 2008, URL: http://www.vattenfall.se/sv/file/2_Technical_Description_Lillgru_8459792.pdf_16611908.pdf [cited 7 December 2011].

[2]Siemens AG, SWT-2.3-93 Product Brochure [published online] Siemens AG Energy Sector, Erlangen, Germany, Siemens Wind Power A/S, Brande, Denmark, 2009, URL: http://www.energy.siemens.com/mx/pool/hq/power-generation/wind-power/E50001-W310-A102-V6-4A00_WS_SWT-2.3-93_US.pdf [cited 7 December 2011].

[3]Dahlberg, J.-Å., "Assessment of the Lillgrund Wind Farm: Power Performance Wake Effects," [online report] Vattenfall Vindkraft AB, 6_1 LG Pilot Report, Sept. 2009, URL: http://www.vattenfall.se/sv/file/15_Assessment_of_the_Lillgrund_W.pdf_16596737.pdf [cited 9 December 2011].

[4]Ivanell, S. S. A., "Numerical Computations of Wind Turbine Wakes," Ph.D. Thesis, Dept. of Mechanics, Gotland Univ., Stockholm, Sweden, 2010.

[5]Mann, J., "The Spatial Structure of Neutral Atmospheric Surface-Layer Turbulence," *Journal of Fluid Mechanics*, Vol. 273, 1994, pp. 141–168.

[6]Calaf, M., Meneveau, C., and Meyers J., "Large Eddy Simulation Study of Fully Developed Wind-Turbine Array Boundary Layers," *Physics of Fluids*, Vol. 22, 2010, p. 015110.

[7]Porté-Agel, F. Wu, Y.-T., and Conzemius, R. J., "Large-Eddy Simulation of Atmospheric Boundary Layer Flow Through Wind Turbines and Wind Farms," *Journal of Wind Engineering and Industrial Aerodynamics*, Vol. 99, No. 4, 2011, pp. 154–168.

[8]Troldborg, N., "Actuator Line Modeling of Wind Turbine Wakes," Ph.D. Thesis, Department of Fluid Mechanics, Technical University of Denmark, Lyngby, Denmark, 2008.

[9]Mikkelsen, R., Sørensen, J. N., Øye, S., and Troldborg, N., "Analysis of Power Enhancement for a Row of Wind Turbines Using Actuator Line Technique," *Journal of Physics: Conference Series*, Vol. 75, 2007, p. 012044.

[10]Troldborg, N., Larsen, G. C., Madsen, H. A., Hansen, K. S., Sørensen, J. N., and Mikkelsen, R., "Numerical Simulations of Wake Interactions Between Two Wind Turbines at Various Inflow Conditions," *Wind Energy*, Published Online at http://onlinelibrary.wiley.com/doi/10.1002/we.433/abstract, 2010.

[11]Moeng, C.-H., "A Large-Eddy Simulation Model for the Study of Planetary Boundary Layer Turbulence," *Journal of the Atmospheric Sciences*, Vol. 41, 1984, pp. 2052–2062.

[12]Smagorinsky, J., "General Circulation Experiments with the Primitive Equations," *Monthly Weather Review*, Vol. 91, 1963, pp. 99–164.

[13]Stoll, R., and Porté-Agel., "Dynamic Subgrid-Scale Models for Momentum and Scalar Fluxes in Large-Eddy Simulations," *Water Resources Research*, Vol. 42, No. W01409, 2006.

[14]Monin, A. S., and Ohukhov A. M., "Basic Laws of Turbulent Mixing in the Surface Layer of the Atmosphere," *Tr. Akad. Nauk SSR Geophiz. Inst.* Vol. 24, No. 151, 1954, pp. 163–187.

[15]Etling, D., Ch. 2 in *Modelling of Atmospheric Flow Fields*, World Scientific, Singapore, 1996, pp. 54–55.

[16]Sørensen, J. N., and Shen, W. Z., "Numerical Modeling of Wind Turbine Wakes," *Journal of Fluids Engineering*, Vol. 124, 2002, pp. 393–399.

[17]Martínez, L. A., Leonardi, S., Churchfield, M. J., and Moriarty, P. J., "A Comparison of Actuator Disc and Actuator Line Wind Turbine Models and Best Practices for Their Use," *50th AIAA Aerospace Sciences Meeting and Exhibit*, Nashville, TN, Jan. 9–12, AIAA, Washington D.C., 2012.

[18]OpenFOAM, Ver. 2.0.x, [available online] Silicon Graphics International Corporation (SGI), URL: http://www.openfoam.org/git.php [accessed 13 December 2011].

[19]Rhie, C. M., and Chow W. L., "Numerical Study of the Turbulent Flow Past an Airfoil with Trailing Edge Separation," *AIAA Journal*, Vol. 21, No. 11, 1983, pp. 1525–1532.

[20]Issa, R. I., "Solution of the Implicitly Discretized Fluid Flow Equations by Operator-Splitting," *Journal of Computational Physics*, Vol. 62, 1985, pp. 40–65.

[21]Chevalier, C., and Pellegrini, F., "PT-Scotch: A Tool for Efficient Parallel Graph Ordering," [online report] *Proceedings of PMAA2006*, Rennes, France, Oct. 2006, URL: http://www.labri.fr/publications/paradis/2006/CP06a [accessed 13 December 2011].

[22]Bergström, H., "Meteorological Conditions at Lillgrund," [online report] Vattenfall Vindkraft AB, 6_2 LG Pilot Report, Mar. 2009, URL: http://www.vattenfall.se/sv/file/16_Meteorological_conditions.pdf_16614584.pdf [cited 7 December 2011].

[23]Charnock, H., "Wind Stress on A Water Surface," Quarterly Journal of the Royal Meteorological Society. Vol. 81, No. 350, 1955, pp. 639–640.

[24]Garratt, J. R., *The Atmospheric Boundary Layer*, Cambridge University Press, Cambridge, 1992, pp. 97–100.

[25]Moeng, C.-H., and Sullivan, P. P., "A Comparison of Shear- and Buoyancy-Driven Planetary Boundary Layer Flows," *Journal of the Atmospheric Sciences*, Vol. 51, 1994, pp. 999–1022.

[26]Laursen, J., Enevoldsen, P., and Hjort, S., "3D CFD Quantification of the Performance of a Multi-Megawatt Wind Turbine," *Journal of Physics: Conference Series*, Vol. 75, No. 012007, 2007.

[27]Leloudas, G., "Optimization of Wind Turbines with Respect to Noise," Master's Thesis, Technical University of Denmark, Lyngby, Denmark, 2006.

[28]Bertagnolio, F., Sørensen, N., Johansen, J., and Fuglsang, P., "Wind Turbine Airfoil Catalogue," Risø National Laboratory, Rept. Risø-R-1280(EN), Roskilde, Denmark, Aug. 2001.

[29]Hansen, C., AirfoilPrep, NWTC Design Code, Ver. 2.0, [available online] National Renewable Energy Laboratory, Golden, CO, 2005, URL: http://wind.nrel.gov/designcodes/preprocessors/airfoilprep/ [last modified 9 March 2010, accessed 7 December 2011].

[30]Du, Z., and Selig, M. S., "A 3-D Stall-Delay Model for Horizontal Axis Wind Turbine Performance Prediction," *36th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV, Jan. 12–15, AIAA, Washington D.C., 1998, pp. 9–19.

[31]Eggers, A. J., Chaney, K., and Digumarthi, R., "An Assessment of Approximate Modeling of Aerodynamic Loads on the UAE Rotor," *41st AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV, Jan. 6–9, 2003, AIAA, Washington D.C., 2003, pp. 283–292.

[32]Viterna, L., and Janetzke, D. C., "Theoretical and Experimental Power from Large Horizontal-Axis Wind Turbines," National Aeronautics and Space Administration, Lewis Research Center, Rept. NASA TM-82944, Cleveland, OH, Sept. 1982.

[33]Jonkman, J., Butterfield, S., Musial, W., and Scott, G., "Definition of a 5-MW Reference Wind Turbine for Offshore System Development," National Renewable Energy Laboratory, Rept. NREL/TP-500-38060, Golden, CO, Jan. 2009.

[34]Khanna, S., and Brasseur, J. G., "Three-Dimensional Buoyancy- and Shear-Induced Local Structure of the Atmospheric Boundary Layer," *Journal of the Atmospheric Sciences*, Vol. 55, 1998, pp. 710–743.

[35]Chamorro, L. P., and Porté-Agel, F., "Flow Characterization of Wind-Turbine Wake(s) Developed in a Boundary Layer Flow with Different Thermal Stratifications: A Wind-Tunnel Study," *The Fifth International Symposium on Computational Wind Engineering*, North Carolina, USA, May 23–27.

[36]Churchfield, M. J., Lee, S., Michalakes, J., and Moriarty, P. J., "A Numerical Study of the Effects of Atmospheric and Wake Turbulence on Wind Turbine Dynamics," submitted *Journal of Turbulence*, 2011.