Large-eddy simulations of idealized atmospheric boundary layers using Nalu-Wind

C M Kaul[1], S Ananthan[2], M J Churchfield[2], J D Mirocha[3], L K Berg[1], R Rai[1]

[1]Pacific Northwest National Laboratory, Richland WA, USA, [2]National Renewable Energy Laboratory, Golden CO, USA [3]Lawrence Livermore National Laboratory, Livermore CA, USA

E-mail: colleen.kaul@pnnl.gov

Abstract. Accurate prediction of wind-plant performance relies, in part, on properly characterizing the turbulent atmospheric boundary layer (ABL) flow in which wind turbines operate. Large-eddy simulation (LES) is a powerful tool for simulating ABLs because it resolves the largest, most energetic scales of three-dimensional turbulent motions. Yet LES predictions are well known to depend on modeling choices such as grid resolution, numerical discretization schemes, and closures for unresolved scales of turbulence. Here, we evaluate how these choices influence predictions of ABL winds using Nalu-Wind, a wind-specific fork of the open-source, generalized, unstructured, massively parallel flow solver NaluCFD/Nalu.

1. Introduction

High-resolution large-eddy simulations (LES) of the turbulent atmospheric boundary layer (ABL) uniquely allow the fully three-dimensional, time-varying operating environment of wind turbines to be characterized [1]. By coupling LES models of the microscale flow to mesoscale atmospheric models, the wide range of flow scales affecting wind-plant performance can be simulated. Understanding the performance of the microscale solver in relevant flow configurations is an important step in tackling the challenges associated with mesoscale–microscale coupling. As part of research conducted under the Mesoscale–Microscale Coupling (MMC) project within the U.S. Department of Energy's (DOE) Atmosphere to Electrons initiative, Mirocha et al. [2] examined the sensitivities of large-eddy simulations of ABL scenarios pertinent to wind energy applications. A suite of simulations was performed using three different code bases exercising a range of options for grid resolutions, advection schemes, and subgrid-scale (SGS) turbulence closures. Forcing parameters of the simulations, such as surface heat fluxes and large-scale winds, were also varied around nominal neutral and convective configurations emulating observed conditions at the Department of Energy/Sandia National Laboratories Scaled Wind Farm Technology (SWiFT) facility in west Texas.

Nalu-Wind [3; 4] is a wind-specific fork of the generalized, unstructured, massively parallel flow solver NaluCFD/Nalu [5] that is designed to enable predictive wind-energy simulations on exascale computing platforms under the Exawind project, part of the DOE's Exascale Computing Project. Towards this goal, it integrates capabilities for simulating aerodynamic interactions with wind turbines and for simulating fine-scale ABL turbulence produced by buoyant forces as well as by wind shear. As the state-of-the-art in high-performance, high

-fidelity modeling for wind energy, Nalu-Wind has been identified as the target microscale model by the MMC project. However, due to its novelty, it was not able to be included in the original MMC comparison study of LES codes. Therefore, this study addresses this gap by using Nalu-Wind to simulate several idealized ABL scenarios. We explore the sensitivity of Nalu-Wind's predictions to variations in model resolution, numerics, and closures, in comparison to observational data and results from other models.

2. Approach

In this section, we describe some key aspects of Nalu-Wind's physical equation set and numerical discretization approaches. Readers seeking a broader discussion of Nalu-Wind's full capabilities may consult the model's documentation [3; 4] and references therein. The models used by Mirocha et al. [2], and whose results are reproduced here, are also briefly outlined. We additionally review the case description and simulation setup provided in Mirocha et al. [2], noting that there are some differences with their discussion for consistency with the actual simulations performed.

2.1. Nalu-Wind Formulation

Nalu-Wind is a generalized, unstructured flow solver that employs the finite-volume discretization approach [3; 4]. Nalu-Wind was forked from the solver NaluCFD/Nalu, which employs a low Mach number approximation to the flow governing equations [5], to facilitate the development of specializations and enhancements for wind energy applications. In particular, Nalu-Wind offers a formulation of the flow governing equations using the Boussinesq approximation, allowing for efficient simulation of ABL flows with vertical depths of several kilometers or less (e.g., [6]). Prognostic equations are solved for the three components of velocity, potential temperature, and, optionally, for the SGS turbulent kinetic energy (TKE).

Two discretizations are available in Nalu-Wind: a control volume finite element method (CVFEM) [7; 8] and an edge-based, vertex-centered (EBVC) scheme that is similar to cell-centered finite volume schemes used in other codes such as OpenFOAM [9]. For highly structured meshes, the EBVC scheme performs well, while for poor quality meshes the CVFEM scheme reduces errors associated with non-orthogonality. The EBVC scheme offers greater computational efficiency, as the evaluation of time and source terms are collapsed to a single quadrature point per nodal volume and, similarly, evaluations of advection and diffusion terms are only required at edge midpoints. This results in a seven point stencil in three dimensions for the EBVC scheme, compared to a 27 point stencil for CVFEM [5]. Our baseline model configuration uses the EBVC scheme, while we employ the CVFEM scheme as a sensitivity test.

Advection stabilization is performed as a weighting between generalized central and higher-order upwind interpolation/extrapolation operators for the advected scalars (velocity component, potential temperature, etc.). The weighting is controlled by a blending function that depends on the cell Peclet number. Two functional forms are available in Nalu-Wind. In most of the simulations presented here, the hyperbolic tangent form blending function is used with transition Peclet number and transition width of 50,000 and 200, respectively, for velocities, and 2 and 1, respectively, for all other scalars. Consequently, in any significantly turbulent region of the flow, velocity values at the integration points are computed with central operators whereas upwind operators are applied to the other scalars. Additionally, a Van Leer limiter is applied to extrapolated values of potential temperature. As a sensitivity test, we perform simulations using pure central operators to compute advective fluxes at integration points.

Nalu-Wind's pressure stabilization method can be briefly described as an incremental fourth-order approximate projection scheme with time-step scaling [8; 10; 11]. Time integration uses an implicit, second-order backward difference formula (BDF2) scheme [12]. All Nalu-Wind simulations presented here have a fixed, 1 s time step unless otherwise noted.

Various subgrid-scale turbulence parameterizations are available in Nalu-Wind. As our baseline configuration, we employ the 1.5-order, SGS TKE closure described in the Nalu-Wind documentation [3]. Note that the mixing length used in this model is purely geometric (i.e. related to the grid spacing) and not limited under local conditions of stable stratification. As a sensitivity test, a standard Smagorinsky closure (again, with no modifications for stratification) is used in its place.

2.2. Other Code Bases

Models used by Mirocha et al. [2], whose results are reproduced here, included the Weather Research and Forecasting (WRF) model [13], the Simulator fOr Wind Farm Applications (SOWFA) [14], and High Gradient (HiGrad) applications model [15]. Among these models, SOWFA is most like Nalu-Wind. Both solve the Boussinesq equation set (the other models solve some form of the compressible equations) and the finite-volume formulation used by SOWFA is similar to the EBVC scheme of Nalu-Wind.

2.3. Case Description and Observations

Neutral and convective cases are defined based on observations from a 200 m instrumented meteorological tower at Texas Tech University's National Wind Institute, adjacent to the Sandia National Laboratories SWiFT test facility [2; 16]

2.4. Simulation Setup

2.4.1. Grid Configuration The filter scale, typically defined as $\Delta = (\Delta x \Delta y \Delta z)^{1/3}$, is commonly used to characterize the resolution of large eddy simulations. Here we are interested in varying the horizontal (Δx, Δy) and vertical (Δz) grid spacings while holding Δ approximately constant. Following Mirocha et al. [2], we therefore consider the effect of varying the aspect ratio $AR = \Delta x/\Delta z$ of each grid cell, with $\Delta y = \Delta x$. For the neutral case, we compare $AR = 3.3$ ($\Delta x = 25$ m, $\Delta z = 7.5$ m) as used by WRF and $AR = 1$ ($\Delta x = \Delta z = 15$ m) as used by HiGrad and SOWFA. For the convective case, we compare $AR = 3$ ($\Delta x = 30$ m, $\Delta z = 10$ m) as used by WRF and $AR = 1$ ($\Delta x = \Delta z = 20$ m) as used by HiGrad.

The neutral case domain extends 2.4 km in each horizontal dimension and 2 km in the vertical. The convective case domain has a 6 km horizontal extent and 3 km vertical depth.

2.4.2. Initial and Reference Conditions The neutral case is initialized in Nalu-Wind by specifying constant initial velocity profiles equal to the geostropic velocity components $(u_g, v_g) = (-2.223\,\mathrm{m\,s^{-1}}, 6.108\,\mathrm{m\,s^{-1}})$, giving a geostrophic wind speed $U_g = 6.5\,\mathrm{m\,s^{-1}}$. The initial profile of potential temperature is given by $\theta = 300$ K at heights below 1 km. Above 1 km, the potential temperature increases at a rate of 10 K km^{-1}. To generate turbulence, potential temperature perturbations, drawn from a standard normal distribution and multiplied by an amplitude factor of 0.25 K, are added below 500 m. To ensure a non-zero turbulent viscosity during the initial spin-up of the simulation, the SGS TKE is initialized with a value of 0.3 $\mathrm{m^2\,s^{-2}}$ at the surface, decreasing linearly to zero by $z = 1100$ m, and zero at all heights above. This surface value is consistent with typical scaling rules of TKE with surface friction velocity u_*.

Initial velocity profiles of the convective case are equated to the geostrophic winds, $(u_g, v_g) = (9\,\mathrm{m\,s^{-1}}, 0\,\mathrm{m\,s^{-1}})$. The mean potential temperature equals 309 K below 600 m, and increases at a rate of 4 K km^{-1} above 600 m. Below 300 m, potential temperature perturbations are added in the same manner as described for the neutral case. SGS TKE is 0.3 $\mathrm{m^2\,s^{-2}}$ at the surface, decreasing to zero at 700 m.

The reference pressure and temperature of the simulations are taken as 1000 hPa and 300 K, respectively, yielding a reference density $\rho = 1.163\,\mathrm{kg\,m^{-3}}$ in the Boussinesq equation system.

2.4.3. Boundary Conditions and Large-Scale Forcings The surface boundary condition for momentum follows the conventional approach in ABL modeling by invoking Monin-Obukhov similarity theory to predict momentum fluxes. A key aspect of this approach is the prescription of empirically determined stability functions. Mirocha et al. [2] unified all participating models to use the same stability function [17]. Here we opt to retain Nalu-Wind's original choice of stability function [18]. This difference affects the convective case, and the results indicate that other differences between the models outweigh sensitivity to this small difference in surface flux specification.

The surface sensible heat flux is fixed in both the convective and neutral cases. The surface potential temperature flux was set to $H_S = 0\,\mathrm{K\,m\,s^{-1}}$ for the neutral case and $H_S = 0.35\,\mathrm{K\,m\,s^{-1}}$ for the baseline convective case of Mirocha et al. [2], equal to a surface sensible heat flux of 408.5 W $\mathrm{m^{-2}}$ in our Nalu-Wind simulations.

For this study, a Rayleigh damping layer was implemented in Nalu-Wind. Velocity fluctuations with respect to the geostrophic velocities are damped in the upper 20% of the vertical domain extent, or above 1600 m in the neutral case and 2400 m in the convective case. Vertical velocity fluctuations are similarly damped towards zero. Also, Nalu-Wind's computation of Coriolis momentum source terms was modified to use an f-plane approximation, rather than the full three-dimensional form, and forcing terms due to the large-scale pressure gradient, assumed to be in balance with the geostropic wind, were added.

3. Results

Here we present the results of our simulations using Nalu-Wind and compare to the observations and results from WRF, SOWFA, and HiGrad models presented in Mirocha et al. [2], following the simulations naming scheme used in that work (i.e. their Tables 1 and 2). As in Mirocha et al. [2], the analysis is performed in the temporal domain using time series of real or virtual tower data. Time series data of observations and simulations with WRF, SOWFA, and HiGrad were retrieved from the Atmosphere to Electrons Data Archive and Portal [19]; statistical analyses are similar to those presented by Mirocha et al. [2], but all quantities were recomputed to ensure consistency with the analysis of Nalu-Wind simulation data. Note that there is a typographical error in Mirocha et al. [2]; in that work, wind speed variability is computed as the variance of the 1 Hz wind speed, but is described in the text as the variance of the ten-minute wind speed. Here, we actually compute the variance of the ten-minute wind speed as a means of characterizing the wind speed variability.

Tables 1 and 2 summarize the simulations performed using Nalu-Wind. The naming convention is based on four attributes of the simulations, as follows: simulations using the "baseline" options (i.e., EBVC discretization, baseline advection stabilization, and SGS TKE closure) are designated by "N" for Nalu-Wind; other simulations are named by the modified option. The second part of the name, "N" or "C" indicates the neutral or convective case. The first number indicates the forcing, "1" for the standard forcing and "2" or "3" for perturbations to the standard forcing. The final number is the grid aspect ratio. An "X" in place of an attribute means we consider all possible options. For example, N-N-X-1 means the set of simulations N-N-1-1, N-N-2-1, and N-N-3-1.

3.1. Neutral Case

We analyze all Nalu-Wind neutral case simulation results over a two-hour window centered at the fifteenth hour of the simulation. This analysis period was determined by examining the time series of the spatially averaged 80 m wind speed for its first peak in value, associated with inertial oscillations; given the magnitude of the wind speed and the domain size, this is equivalent to an averaging window of several minutes, consistent with the approach taken by Mirocha et al. [2]. As the average wind speed shows rather small variations for several hours around the timing of

Table 1. Summary of neutral case simulations. The baseline options are described in Section 2.1.

Run name	Δx (m)	Δz (m)	U_g (m s^{-1})	z_0 (m)	Comment
N-N-1-3.3	25	7.5	6.5	0.05	Baseline options
N-N-1-1	15	15	6.5	0.05	Baseline options
N-N-2-3.3	25	7.5	7.15	0.1	Forcing sensitivity
N-N-2-1	15	15	7.15	0.1	Forcing sensitivity
N-N-3-3.3	25	7.5	5.85	0.01	Forcing sensitivity
N-N-3-3.3	15	155	5.85	0.01	Forcing Sensitivity
CVFEM-N-1-3.3	25	7.5	6.5	0.05	uses CVFEM scheme
CVFEM-N-1-1	15	15	6.5	0.05	uses CVFEM scheme
CentAd-N-1-3.3	25	7.5	6.5	0.05	central advection interpolation
CentAd-N-1-1	15	15	6.5	0.05	central advection interpolation
Smag-N-1-3.3	25	7.5	6.5	0.05	uses Smagorinsky SGS closure
Smag-N-1-1	15	15	6.5	0.05	uses Smagorinsky SGS closure

Table 2. Summary of convective case simulations. The baseline options are described in Section 2.1.

Run name	Δx (m)	Δz (m)	U_g (m s^{-1})	H_s (K m s^{-1})	Comment
N-C-1-3	30	10	9.0	0.3500	Baseline options
N-C-1-1	20	20	9.0	0.3500	Baseline options
N-C-2-3	30	10	10.0	0.4364	Forcing sensitivity
N-C-2-1	20	20	10.0	0.4364	Forcing sensitivity
CVFEM-C-1-3	30	10	9.0	0.3500	uses CVFEM scheme
CVFEM-C-1-1	20	20	9.0	0.3500	uses CVFEM scheme
CentAd-C-1-3	30	10	9.0	0.3500	central advection interpolation
CentAd-C-1-1	20	20	9.0	0.3500	central advection interpolation
Smag-C-1-3	30	10	9.0	0.3500	uses Smagorinsky SGS closure
Smag-C-1-1	20	20	9.0	0.3500	uses Smagorinsky SGS closure

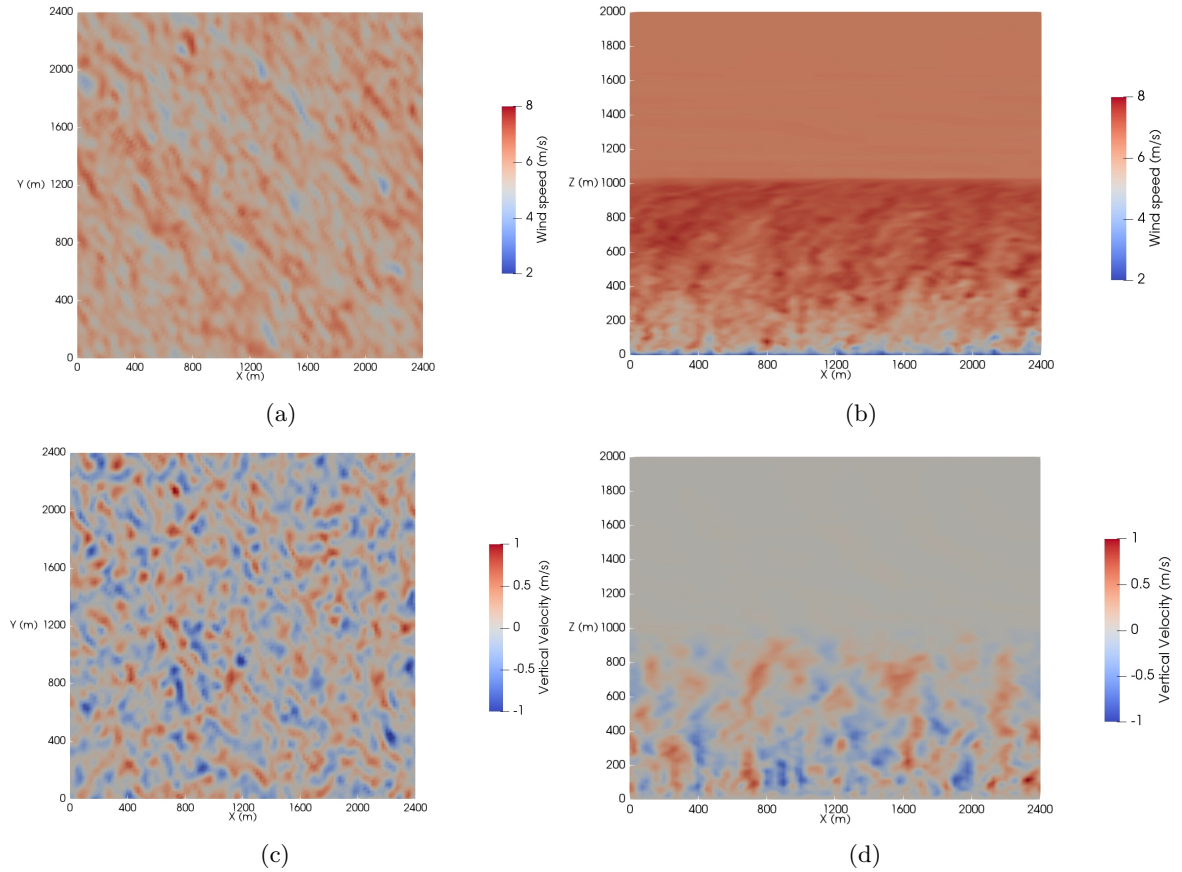

Figure 1. Instantaneous snapshots at 14 hours simulated time of baseline neutral case with $AR = 3.3$ (N-N-1-3.3). The horizontal wind speed (top row) and vertical velocity (bottom row) are shown on (a,c) a horizontal plane at 100 m; (b,d) on an east–west plane through the center of the domain.

the peak, the results should not be critically sensitive to the precise specification of the analysis window.

3.1.1. Baseline Configurations Instantaneous snapshots at 14 hours simulated time of the baseline neutral simulation with $AR = 3.3$ (designated as simulation N-N-1-3.3 in Table 1) are shown in Figure 1. Examining the wind speed on a plane at 100 m (panel a) shows elongated structures oriented with the mean flow direction, while the cross-stream vertical slice (panel b) clearly shows the deceleration of the flow near the surface as well as the effectiveness with which the stable potential temperature gradient above 1000 m limits the vertical growth of the turbulent ABL. Vertical velocity fluctuations (panels c and d) are smaller in magnitude and more disorganized in structure.

Figure 2 compares wind speed U, normalized by the friction velocity u_* predicted by each model over the lowest quarter of the ABL (i.e. lowest 250 m). Results with WRF, SOWFA, and HiGrad are from the W1, S1, and H1 simulations defined by Mirocha et al [2], respectively, as these have the greatest similarity in terms of grid configuration, forcing, and type of subgrid-scale turbulence model to the Nalu-Wind simulations. Theory predicts a logarithmic variation of Uu_*^{-1} with height z from the surface. We see that all models tend to predict faster increase of wind speed with height than the theoretical rate and that none stand out as especially

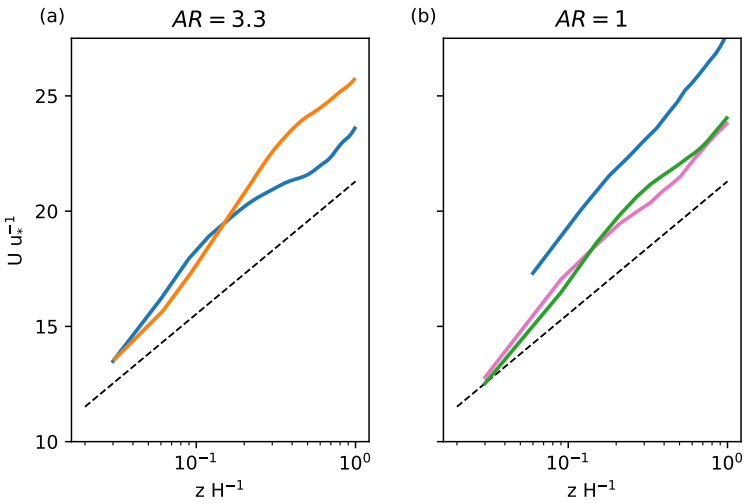

Figure 2. Comparison of mean wind speed U, scaled by the surface friction velocity u_*, to the theoretical log-law scaling with von Kármán constant $\kappa = 0.4$ (black dashed line) for each model. Height scale $H = 250$ m and roughness height $z_0 = 0.05$ m. Panel (a) shows results using a grid aspect ratio of 3.3 for Nalu-Wind (N-N-1-3.3, blue) and WRF (W1, orange). Panel (b) shows results using a grid aspect ratio of 1 for Nalu-Wind (N-N-1-1, blue), SOWFA (S1, pink), and HiGrad (H1, green).

inferior or superior to the others. Comparing the N-N-1-3.3 and N-N-1-1 simulations illustrates the sensitivity of the results to the numerical grid, as these simulations differ only in this regard. In particular, while N-N-1-1 shows an approximately logarithmic variation of wind speed with height (albeit with the wrong slope), the wind speed values of N-N-1-3.3 do not vary logarithmically, first increasing too fast up to about 50 m and then increasing too slowly. Note that in the version of Nalu-Wind tested here, the height above the surface is taken to be one quarter of the length of the nearest edge which intersects the boundary face. The implementation of the surface scheme in newer versions of Nalu-Wind will allow the user to specify the height at which to evaluate the similarity functions.

Another perspective on predicted mean wind speeds is provided by Figure 3, which plots means of the ten-minute running average wind speed over the two hour analysis window (less the first ten minutes) for each model and observations. Shading or bars are used to indicate a range of one standard deviation σ of the ten-minute average wind speed from the overall mean. The models show a low level of variability relative to the observations, with Nalu-Wind and SOWFA having especially low variability. This discrepancy between simulations and observations is possibly attributable on one hand to the lack of variability in large-scale forcing in the idealized simulations (e.g., constant values of U_g) and on the other, to tower-wake artifacts in the observations [2]. The limited domain size of the simulations may also play a role. Note that since we are comparing ten-minute averaged wind speeds, the grid-scale filtering implicit in the simulations, which can be associated with length scales of tens of meters and time scales of several seconds, should not dominate the comparison of variability. Note that although HiGrad and SOWFA showed similar scaled velocity profiles in Figure 2, they show significant differences in the wind speed profiles plotted here. Conversely, although Nalu-Wind and SOWFA do not agree well in Figure 2, the velocity profiles themselves match closely.

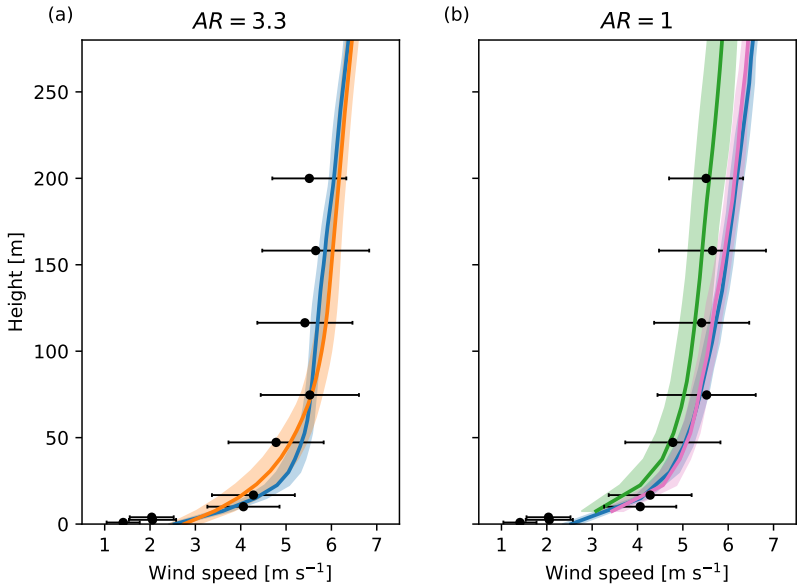

Figure 3. Mean wind speeds in the baseline neutral case for (a) Nalu-Wind (N-N-1-3.3, blue) and WRF (W1, orange) and (b) Nalu-Wind (N-N-1-1; blue), SOWFA (S1, pink) and HiGrad (H1, green). Ten-minute average wind speeds U_{10min} are computed over a two-hour period. The averages of U_{10min} are shown as solid lines; their standard deviations are shown by the shaded regions. Means of the observed wind speeds are plotted as dots, while the bars show the standard deviations.

3.1.2. Sensitivity to Forcing Parameters In this section, we consider model response to changes in simulation forcing parameters. Specifically, we vary two important aspects of the forcing, the roughness length z_0 and the geostrophic wind speed U_g, to produce two sets of sensitivity cases, N-N-2-X and N-N-3-X, as listed in Table 1. These can be compared to the sets of simulations W2, S2, H2 and W3, S3, H3 described in Mirocha et al. [2]. Note that to reduce the total number of simulations, z_0 and U_g are varied together, not independently.

Figure 4 shows the effect of the forcing perturbation on the scaling of the normalized wind speed. The dashed reference lines are recomputed in each panel to capture the expected change in slope depending on the roughness length. Nalu-Wind simulations show similar behavior to the baseline case: the $AR = 1$ simulations show a logarithmic trend albeit with too great a slope, while the $AR = 3.3$ simulations first overpredict, then underpredict, the increase of wind speed with height. These errors are slightly worsened with increasing z_0. One interesting feature of the simulations is the persistently greater values of Uu_*^{-1} in Nalu-Wind simulations with $AR = 1$ (N-N-X-1 cases). This difference is mostly attributable to the systematically lower values of u_* predicted in the $AR = 1$ Nalu-Wind simulations rather than to differences in the mean windspeed, as the $AR = 3.3$ simulations predict higher U below 100 m for all three forcing specifications.

The sensitivity of the mean wind speed profile to increasing U_g and z_0 is shown in Figure 5, while the response to decreasing U_g and z_0 is shown in Figure 6. The results shown as dotted lines correspond to the baseline case (i.e., profiles plotted in Figure 3). Above the lowest few grid levels, the response of mean wind speed U is dominated by the change in U_g and follows the expected trend: greater U_g drives greater U. However, the models differ in their sensivitity. For example, in Figure 5, the change in hub-height wind speed predicted by WRF is nearly

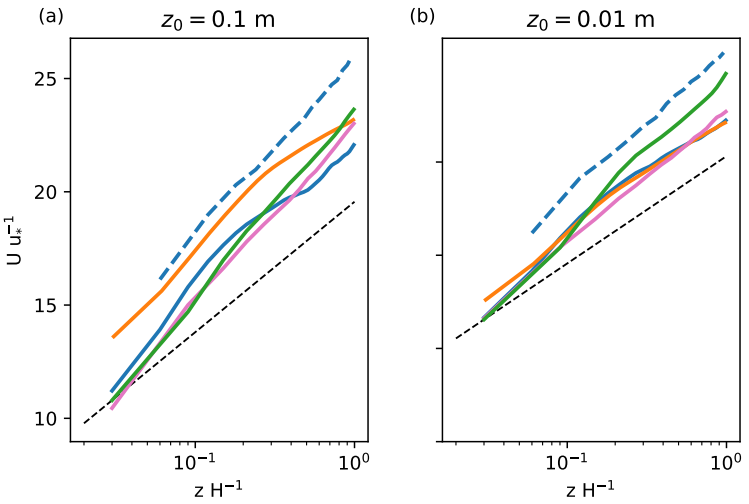

Figure 4. As in Figure 2, showing the effects of modifying U_g and z_0. Results are shown in (a) for $U_g = 1.1U_{g,0}$, $z_0 = 0.1$ m for Nalu-Wind, $AR = 3.3$ (N-N-2-3.3, solid blue); Nalu-Wind, $AR = 1$ (N-N-2-1, dashed blue); WRF (W2, orange); SOWFA (S2, pink); and HiGrad (H2; green) and (b) for $U_g = 0.9U_{g,0}$, $z_0 = 0.01$ m for Nalu-Wind, $AR = 3.3$ (N-N-3-3.3, solid blue); Nalu-Wind, $AR = 1$ (N-N-3-1, dashed blue); WRF (W3, orange); SOWFA (S3, pink); and HiGrad (H3, green).

one-to-one with the change in U_g, while HiGrad shows almost no change. The response of Nalu-Wind and SOWFA is intermediate. Setting aside HiGrad, we see that differences between the models are amplified as U_g is increased. This is significant because of the highly non-linear relationship between wind speed and power: uncertainties in wind speed map in a complex way onto uncertainties in wind plant performance metrics.

3.1.3. Sensitivity to Model Configuration Finally, in this section, we focus on results obtained using Nalu-Wind, varying features of the numerical discretization and subgrid scale closures. These sensitivity cases are listed in Table 1. Figure 7 compares the velocity scaling. We see similar responses among the two groups of simulations using $AR = 3.3$ and $AR = 1$. In particular, the response to changing the treatment of advection stabilization (CentAd-N-1-X) is the weakest. This indicates that results for this neutral case have limited sensitivity to the details of the scalar advection treatment, recalling that central schemes are applied for the velocity advection regardless. The strongest response results from changing the subgrid-scale model (Smag-N-1-X). An intermediate level of sensitivity is shown to the use of the CVFEM discretization approach (CVFEM-N-1-X). As might be expected, use of the Smagorinsky closure without modifications for wall-bounded flow yields poor agreement with the theoretical scaling.

Mean wind speed profiles are compared in Figure 8. When not normalizing by the friction velocity, trends among the simulations are less distinct.Additionally, note that we have not attempted to optimize SGS closure coefficients for this case, and smaller differences between the simulations using Smagorinsky and TKE-based closures might be attainable by tuning coefficient values. Here, the grey shading indicates the maximum plus or minus one standard deviation range in ten minute average wind speed among all the simulations with a given grid aspect ratio. Remarkably, the range of variability remains smaller than the variability of the observations at all measurement heights. Additionally, sensitivity to the model configuration depends on the grid configuration and the specific heights at which wind speed data is needed.

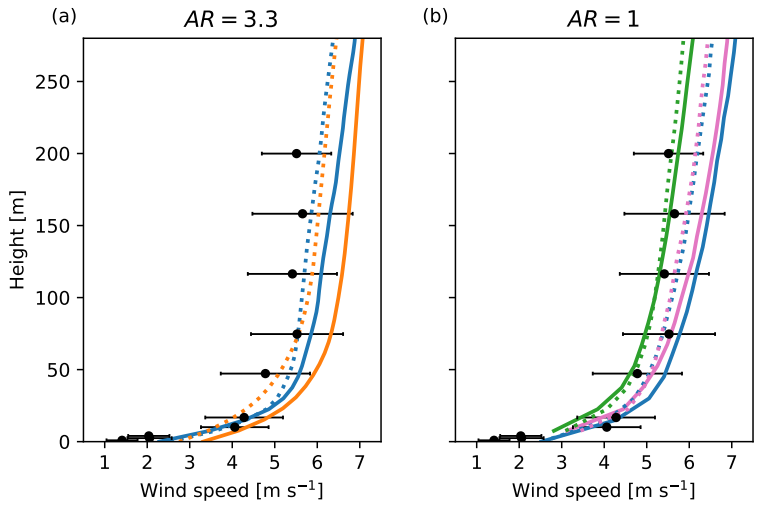

Figure 5. Effect of increasing U_g and z_0 on mean windspeed (computed as described for Figure 3). Results plotted with solid lines have $U_g = 1.1U_{g,0}$, $z_0 = 0.1$ m (run "2" for each model), while dotted lines show the baseline case results for the corresponding model and grid configuration. Panel (a) shows Nalu-Wind, $AR = 3.3$ (blue) and WRF (orange). Panel (b) shows Nalu-Wind, $AR = 1$ (blue), SOWFA (pink), and HiGrad (green). Observations are shown as in Figure 3.

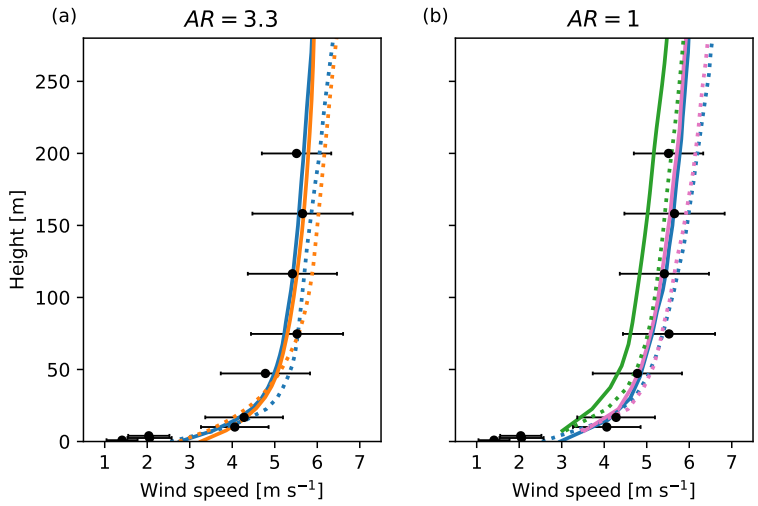

Figure 6. As in Figure 5, but for $U_g = 0.9U_{g,0}$, $z_0 = 0.01$ m (run "3" for each model).

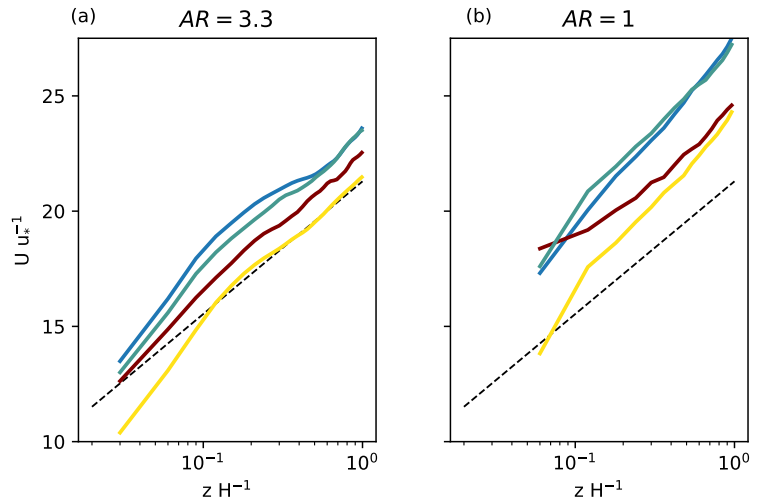

Figure 7. Sensitivity of mean wind speed scaling (plotted as in Figure 2) to solution options in Nalu-Wind: baseline options (N-N-1-3.3 and N-N-1-1, blue); CVFEM discretization (CVFEM-N-1-3.3 and CVFEM-N-1-1, maroon); central scheme for interpolation of advected scalars (CentAd-N-1-3.3 and CentAd-N-1-1, teal); and Smagorinsky SGS closure (Smag-N-1-3 and Smag-N-1-1, yellow). $AR = 3.3$ simulations are shown in panel (a) and $AR = 1$ simulations are shown in panel (b).

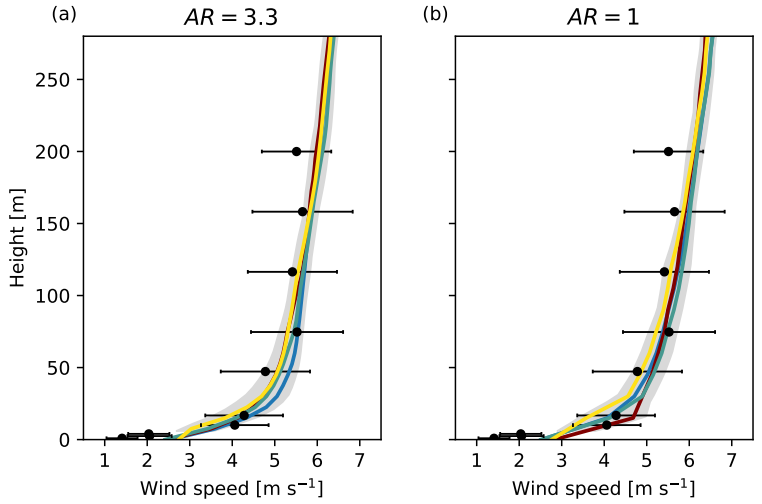

Figure 8. Sensitivity of mean wind speed (computed as in Figure 3) to solution options in Nalu-Wind. Line styles of Nalu-Wind solution results follow the conventions of Figure 7. $AR = 3.3$ simulations are shown in panel (a) and $AR = 1$ simulations are shown in panel (b). Gray shading is the maximum plus/minus one standard deviation range of ten-minute wind speeds among all simulations plotted in a panel. Means and variability of observed winds are plotted as in Figure 3.

3.2. Convective Case
Convective case simulations are analyzed similarly to those of the neutral case. Due to the stronger turbulence, a shorter simulation spin up time of only one hour is required, and following Mirocha et al. [2] we composite statistics over the following two hours.

3.2.1. Baseline Configurations Figure 9 shows snapshots of the horizontal wind speed and vertical velocity in the $AR = 3$ baseline configuration simulation (N-C-1-3). The change in atmospheric stability regime from the neutral case (Figure 1) is obvious. Fluctuations in wind speed and vertical velocity are much larger and are correlated over the ABL depth: this is particularly visible in the coherent updraft structures shown in panel d. Quasi-cellular, rather than streaky, flow structures are apparent in the horizontal views (panels a and c).

Figure 10 shows the mean wind predicted by Nalu-Wind in the baseline configuration, in comparison to observations and results of the most comparable simulations using the WRF and HiGrad models (W1 and H1). The analysis follows that described in Figure 3. In particular the shading or bars show plus and minus one standard deviation of the ten-minute average wind speed. In contrast to the neutral case, the variability of the simulations is comparable to, or even exceeds, the observed variability. Here, it would be informative to analyze the variance of the vertical velocity and how the strength of the simulated convective updrafts affects fluctuations of the horizontal winds. For the Nalu-Wind simulations, the sensitivity of the mean wind speed to the grid cell aspect ratio is small, particularly in relation to the temporal variability of the results. For all models, simulated wind speeds are biased low relative to the observations, but above approximately 50 m, all modeled variability ranges overlap with the observational range.

3.2.2. Sensitivity to Forcing Parameters In this section, we examine the sensitivity of the results to increasing the surface heat flux (linearly related to the surface potential temperature flux denoted H_S) and the geostrophic velocity U_g. A 25 percent increase in H_S and 1 m s^{-1} increase in U_g are applied. The resulting mean wind speeds and one standard deviation range of ten-minute averaged wind speeds are shown in Figure 11. The $AR = 3$ simulations with Nalu-Wind and WRF have a similar response, each yielding a nearly one-to-one ratio of changes in hub-height wind speed and U_g. In agreement with our expectation that higher surface heat fluxes would drive stronger turbulence, the perturbed Nalu-Wind simulation (N-C-2-3) shows significantly greater variability of ten-minute average wind speeds relative to the baseline (N-C-1-3). Wind speed increases more weakly in the perturbed, $AR = 1$ simulation using Nalu-Wind (N-C-2-1). Curiously, the variability of the wind speed does not change significantly. To examine whether this finding was a peculiarity of the ten-minute wind speed data, or more fundamental, we also computed the resolved turbulent kinetic energy of the simulations (not shown). A much weaker increase in resolved TKE is found between the N-C-1-1 and N-C-2-1 simulations than for the other simulation pairs.

3.2.3. Sensitivity to Model Configuration In Figure 12 we assess the sensitivity of predicted mean wind speed to the same set of modifications applied in the neutral case: the use of the CVFEM discretization, instead of the edge-based scheme; disabling of advection stabilization; and use of a Smagorinsky subgrid-scale turbulence closure, instead of a 1.5-order, TKE-based closure. Among the $AR = 3$ simulations, the greatest sensitivity is due to the use of the Smagorinsky closure (Smag-C-1-3) below about 50 m. Overall, the differences between the mean wind speed predictions are fairly small relative to the large variability of the ten-minute average wind speed during the analysis period. The $AR = 1$ simulations behave differently. While the sensitivity to changing the SGS closure is comparable to the sensitivity found in the $AR = 3$ simulations, turning off advection stabilization for scalars has a much more significant effect.

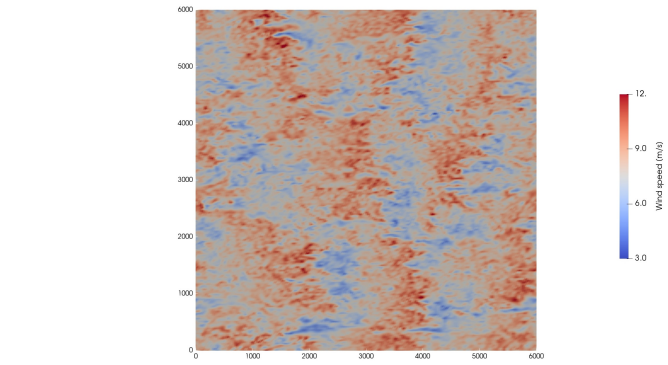

(a)

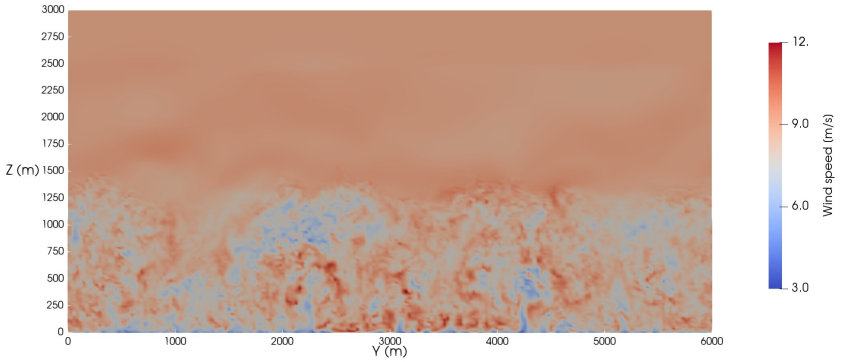

(b)

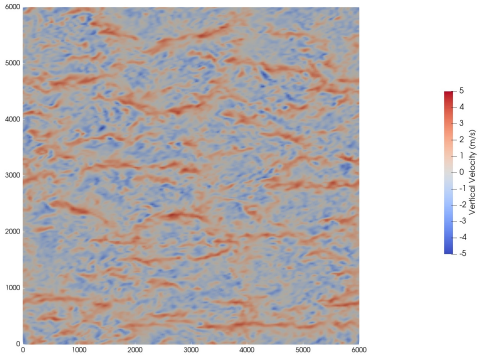

(c)

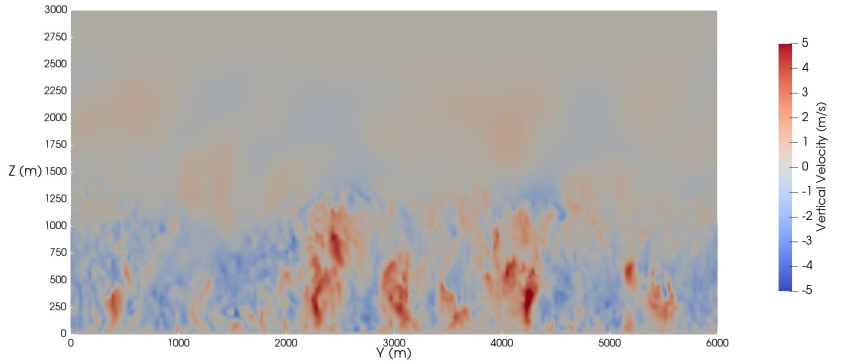

(d)

Figure 9. Instantaneous snapshots at 1 hour simulated time of baseline convective case with $AR = 3$ (N-C-1-3). The horizontal wind speed (a,b) and vertical velocity (c,d) are shown on a horizontal plane at 100 m (a,c); on an east–west plane through the center of the domain (b,d).

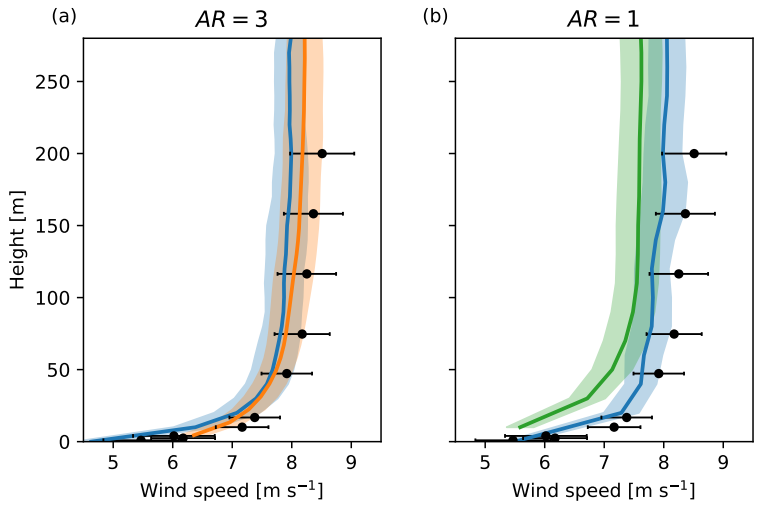

Figure 10. Mean wind speeds in the baseline convective case for (a) Nalu-Wind (N-C-1-3, blue) and WRF (W1, orange) and (b) Nalu-Wind (N-C-1-1; blue) and HiGrad (H1, green). Ten-minute average wind speeds U_{10min} are computed over a two-hour period. The averages of U_{10min} are shown as solid lines; their standard deviations are shown by the shaded regions. Means of the observed wind speeds are plotted as dots, while the bars show the standard deviations.

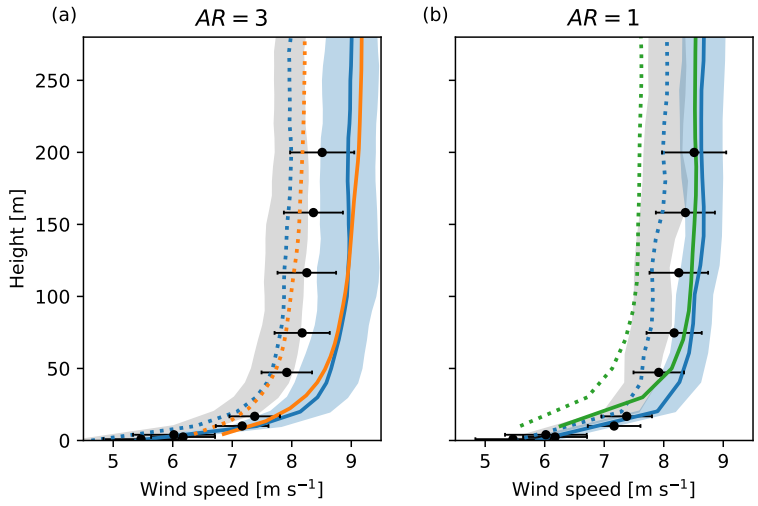

Figure 11. Effect of increasing U_g and surface flux of θ, H_S on mean wind speed (computed as described for Figure 10). Results plotted with solid lines have $U_g = 10$ m s^{-1}, $H_S = 1.25H_{S,0}$ m (run W4 and H3 for WRF and HiGrad, respectively, and runs N-C-2-3 and N-C-2-1 for Nalu-Wind), while dotted lines show the baseline case results for the corresponding model and grid configuration. Panel (a) shows Nalu-Wind, $AR = 3$ (blue) and WRF (orange). Panel (b) shows Nalu-Wind, $AR = 1$ (blue) and HiGrad (green). Observations are shown as in Figure 10. The blue and gray shaded areas are $\pm\sigma$ ranges of ten-minute averaged wind speed for the perturbed and baseline forcings of Nalu-Wind, respectively.

Most notably, the mean wind speed profile of the CVFEM sensitivity simulation (CVFEM-C-1-1, the solid line) is not reasonably smooth. To verify the jagged profile is not an artifact of interpolation to the virtual tower location, we also plotted instantaneous, domain-averaged wind speed profiles and found the same issue. A check of maximum Courant numbers occurring during the simulation did show significantly higher values occurring for CVFEM-C-1-1 than for the baseline simulation, N-C-1-1. Therefore, we performed a second simulation of the CVFEM-C-1-1 configuration, halving the time step to $\Delta t = 0.5$ s. The resulting maximum Courant number at any point in the simulation did not exceed 0.53. The unevenness of the profile is reduced, but not eliminated. Notably, the mean hub-height wind speed is shifted about 0.5 m s^{-1} slower using the reduced time step.

To examine the question of time step sensitivity, we performed another two simulations of the CVFEM-C-1-1 case. One of these runs used a further reduced time step of $\Delta t = 0.25$ s, while the other used $\Delta t = 1$ s but doubled the number of outer iterations used in the solution algorithm. We also performed another run of the N-C-1-1 case using a reduced time step $\Delta t = 0.25$ s. These results are shown in Figure 13. Panel (a) indicates a large sensitivity to timestep or iteration count for the CVFEM simulations and a smaller, but non-negligible sensitivity for the edge-based simulations. However, panel (b) shows the profiles for each discretization type collapse when averaged over the entire periodic flow domain, as well as in time. In other words, the pointwise wind speed statistics are not sufficiently converged over the two hour analysis window due to the high variance of the wind speed. In light of this finding, we re-evaluated the results shown in Figure 10 using the combined planar and temporal averaging. Key features of the model configuration sensitivity are qualitatively in agreement to the sensitivity shown in the time-averaged only results, but the quantitative spread in the models is affected.

4. Conclusions

This study compares idealized simulations of neutral and convective ABLs performed with Nalu-Wind to observations and simulations with other LES models presented by Mirocha et al. [2]. While differences between the models and observations are expected due to the approximate, idealized nature of the simulation set up and forcing, the differences among the WRF, HiGrad, SOWFA, and Nalu-Wind results are difficult to explain due to the widely varying natures of these code bases. Nalu-Wind generally exhibited performance that was plausible relative to the observations and other model results, although there are indications that the implementation of the surface flux scheme and modeling of near-wall turbulence should be improved. Sensitivities of Nalu-Wind to forcing perturbations as well as various discretization and turbulence modeling options were examined. Over all, the strongest sensitivities resulted from forcing perturbations. This strongly argues for the importance of providing realistic forcing of large eddy simulations, using techniques such as those being developed under the DOE Mesoscale–Microscale Coupling project.

Acknowledgments

This research was supported by the Wind Energy Technologies Office of the US Department of Energy Office of Energy Efficiency and Renewable Energy. The Pacific Northwest National Laboratory is operated for the DOE by Battelle Memorial Institute under Contract DE-AC05-76RLO1830.

The research was performed using computational resources sponsored by the Department of Energy's Office of Energy Efficiency and Renewable Energy and located at the National Renewable Energy Laboratory.

LLNL's contribution prepared under Contract DE-AC52-07NA27344.

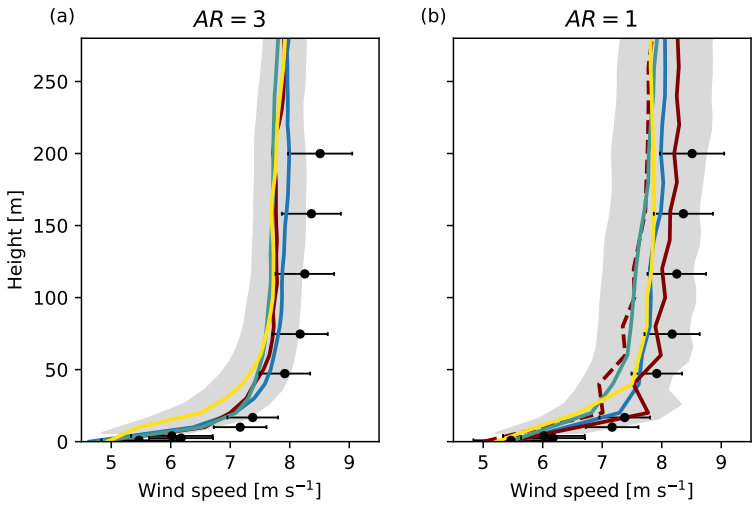

Figure 12. Sensitivity of mean wind speed to solution options in Nalu-Wind: baseline options (N-C-1-3 and N-C-1-1, blue); CVFEM discretization (CVFEM-C-1-3 and CVFEM-C-1-1; maroon); central operators for interpolation of advected scalars (CentAd-C-1-3 and CentAd-C-1-1,teal); and Smagorinsky SGS closure (Smag-C-1-3 and Smag-C-1-1, yellow). $AR = 3$ simulations are shown in panel (a) and $AR = 1$ simulations are shown in panel (b). In panel (b), results are shown for two different runs of the CVFEM-C-1-1 set up. The results shown as a solid line are for a simulation using $\Delta t = 1$ s, while those shown as a dashed line are for a simulation using $\Delta t = 0.5$ s. Gray shading is the maximum plus/minus one standard deviation range of ten-minute wind speeds among all simulations plotted in a panel. Means and variability of observed winds are plotted as in Figure 10.

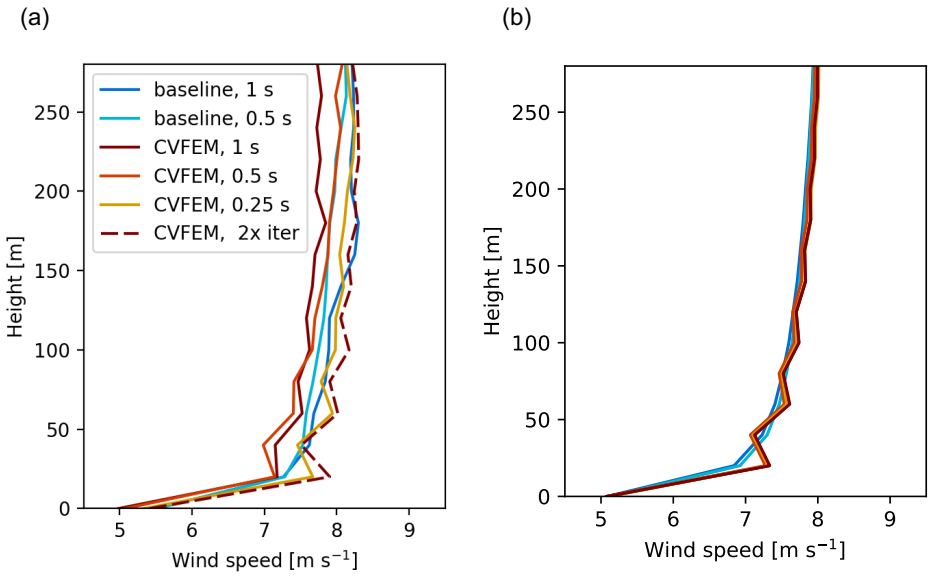

Figure 13. Sensitivity of mean wind speed to time step plotted as (a) time average of single point data and (b) time and planar averages over entire periodic flow domain.

References

[1] Stevens R J A M and Meneveau C 2017 *Annual Review of Fluid Mechanics* **49** 311–339

[2] Mirocha J D, Churchfield M J, Muñoz Esparza D, Rai R K, Feng Y, Kosović B, Haupt S E, Brown B, Ennis B L, Draxl C, Sanz Rodrigo J, Shaw W J, Berg L K, Moriarty P J, Linn R R, Kotamarthi V R, Balakrishnan R, Cline J W, Robinson M C and Ananthan S 2018 *Wind Energy Science* **3** 589–613

[3] Nalu-wind documentation `https://nalu-wind.readthedocs.io/en/latest/index.html`

[4] Nalu-wind Github repository `https://github.com/Exawind/nalu-wind`

[5] Domino S `https://githubcom/NaluCFD/NaluDoc`, 2015 Sierra Low Mach Module: Nalu Theory Manual 1.0

[6] Wyngaard J C 2010 *Turbulence in the Atmosphere* (Cambridge: Cambridge University Press)

[7] Schneider G E and Raw M J 1987 *Numerical Heat Transfer* **11** 363–390

[8] Domino S 2006 *Center for Turbulence Research Summer Proceedings*

[9] OpenCFD 2019 *OpenFOAM – The Open Source CFD Toolbox, User's Manual, Version v1906*

[10] Domino S 2008 *Center for Turbulence Research Summer Proceedings*

[11] Domino S 2014 *Center for Turbulence Research Summer Proceedings*

[12] Ober C C and Shadid J N 2004 *Journal of Computational Physics* **195** 743–772

[13] Skamarock W and Co-authors *A Description of the Advanced Research WRF Version 3, NCAR Technical Note NCAR/TN-475+STR*

[14] Churchfield M J, Lee S, Michalakes J and Moriarty P J 2012 *Journal of Turbulence*

[15] Sauer J A, Muñoz-Esparza D, Canfield J M, Costigan K R, Linn R R and Kim Y J 2016 *Journal of the Atmospheric Sciences* **73** 2615–2632

[16] Kelley C L and Ennis B L 2016 SWiFT Site Atmospheric Characterization, SAND-2016-0216 Tech. rep. Sandia National Laboratories

[17] Stull R B 1988 *An Introduction to Boundary Layer Meteorology* (Kluwer Academic Publisher)

[18] Dyer A J and Hicks B B 1970 *Quarterly Journal of the Royal Meteorological Society* **96** 715–721

[19] Atmosphere to Electrons Data Archive and Portal, `https://a2e.energy.gov/data`