



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Mesoscale and Large-Eddy Simulations for Wind Energy

N. Marjanovic

March 2, 2011

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Mesoscale and Large-Eddy Simulations for Wind Energy

NIKOLA MARJANOVIC, FOTINI KATOPODES Chow, University of California, Berkeley, CA, USA

LLNL-SR-471624

Abstract

Operational wind power forecasting, turbine micrositing, and turbine design require high-resolution simulations of atmospheric flow over complex terrain. The use of both Reynolds-Averaged Navier Stokes (RANS) and large-eddy (LES) simulations is explored for wind energy applications using the Weather Research and Forecasting (WRF) model. To adequately resolve terrain and turbulence in the atmospheric boundary layer, grid nesting is used to refine the grid from mesoscale to finer scales. This paper examines the performance of the grid nesting configuration, turbulence closures, and resolution (up to as fine as 100 m horizontal spacing) for simulations of synoptically and locally driven wind ramping events at a West Coast North American wind farm. Interestingly, little improvement is found when using higher resolution simulations or better resolved turbulence closures in comparison to observation data available for this particular site. This is true for week-long simulations as well, where finer resolution runs show only small changes in the distribution of wind speeds or turbulence intensities. It appears that the relatively simple topography of this site is adequately resolved by all model grids (even as coarse as 2.7 km) so that all resolutions are able to model the physics at similar accuracy. The accuracy of the results is shown in this paper to be more dependent on the parameterization of the land-surface characteristics such as soil moisture rather than on grid resolution.

Introduction

Wind power forecasting is plagued with difficulties in accurately predicting the occurrence and intensity of atmospheric conditions at the heights spanned by industrial-scale turbines (~ 40 -200m). Better simulation of the relevant physics would enable operational practices such as integration of large fractions of wind power into power grids, scheduling maintenance on wind energy facilities, and defining design criteria for next-generation turbines and siting. Increases in available computational power have made high-resolution simulations of the atmospheric boundary layer more practical. The model operator, however, is required to make proper choices regarding model grid spacing, turbulence parameterization, land surface representations and the configuration of initial and lateral boundary conditions for accurate simulations. The “proper” choice of these parameters is highly dependent on the situation under consideration.

This paper investigates the steps necessary to achieve accurate simulations of the wind fields and “ramping” events at a wind farm on the West Coast of the United States. “Ramping” refers to a rapid increase in wind speed over a short time period and is of great interest to wind farm operators for preventing turbine fatigue and predicting power output. The exact definition of a ramping event varies in practice; examples include a 20% capacity change in production over a thirty-minute period [3], or a pre-specified change in the magnitude of system-wide production.

Atmospheric conditions have a large effect on the predictability of a ramping event. Meteorologically-driven ramping events may be due to atmospheric motions at several scales. For example, thermally-driven winds such as sea breezes, mountain-valley circulations, or low-level jets occur at local scales, whereas the passage of a cold front is a synoptically driven phenomenon. In all these cases, associated wind transitions simply appear as large ramps in the wind speed time series to wind farm systems operators, though the cause for the wind changes can be quite different [3]. Larger, synoptically driven features have longer time scales and should in theory be more straightforward to forecast than local-scale phenomena, which usually require fine scale information about land-surface conditions and turbulent mixing in the atmosphere. Low resolution mesoscale models may also be adequate for many situations if the major features of the terrain are adequately captured.

The wind farm which is the focus of this study covers an area of about 10 km by 10 km with a fairly flat topography that includes a few hills with elevation changes less than 150 m (see figure 1). Observation data is available from a sodar, and cup anemometers on the turbine nacelles that can be used for comparison to numerical simulations. We used the Weather Research and Forecasting (WRF) model as our simulation tool [9]. WRF is intended mainly for mesoscale atmospheric simulations and includes large-eddy simulation (LES) capabilities. WRF has a fully compressible, Eulerian and non-hydrostatic equation set. This work focuses on the numerical aspects of the simulation needed to accurately represent the flow over the wind farm.

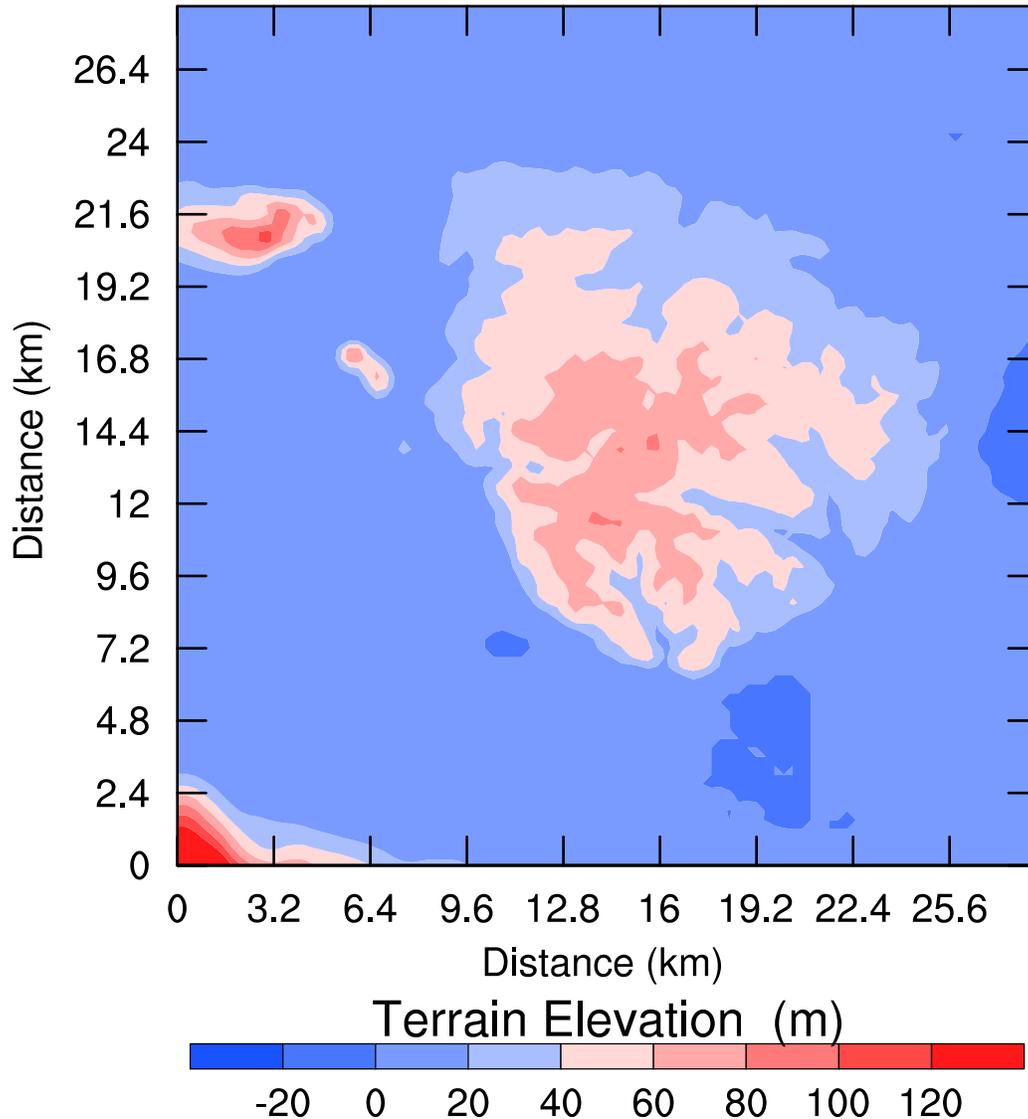


Figure 1: Contours of wind farm topography

Here we examine two ramping events to develop preliminary insight into the ability of WRF to capture shifts in wind speed that are important to wind farm operators. One of these events falls in the category of synoptically-driven transitions and the other in locally-driven transitions. Many studies over complex terrain

point to increasing grid resolution as a means to achieving better agreement of simulations with observations (see e.g. [5]; [4]), and this should be more important for locally-driven transitions. Errors in wind speed can also be due to errors in the representation of turbulent motions, as well as to sub-grid features in the topography and land use, as seen in the simulations of Hanna and Yang (2001). The effects of horizontal and vertical grid resolution, one-way vs. two-way nesting, turbulence closure models, and long term simulation results are examined here. Results are presented below, along with discussion of the numerical setup and grid nesting approaches.

Numerical Simulation Setup

Grid Nesting and topography

Five one-way nested grids were used to simulate flow conditions for the wind farm at horizontal resolutions of 8.1 km, 2.7 km, 900 m, 300 m, and 100 m. The main features of the topography are visible at 2.7 km, but fine structures become much better resolved at 100 m resolution. Topography was extracted from a 10 m resolution data set available from the USGS. The terrain is smoothed near the boundary for each nested sub-domain to match the elevations for the surrounding coarser grid. WRF uses a terrain-following, hydrostatic-pressure vertical coordinate system with the top of the model being a constant pressure surface. The vertical grid spacing is normally assigned by default by WRF, or it can be specified by the user to decrease spacing near the bottom of the model and stretch it to the top. The minimum vertical grid spacing (Δz_{min}) at the surface, as well as the average spacing (Δz_{avg}), are listed in Table 1 for the grid configurations used. The domain height is approximately 13 km in all cases. As seen in table 1, the horizontal spacing (Δh) is uniform and increases in spatial resolution require decreases in the model time-step.

(nx,ny,nz)	Δh (m)	Δz_{min} (m)	Δz_{avg} (m)	Δt (s)
(96,96,49)	8100	35	246	30
(96,96,49)	2700	35	246	10
(96,96,49)	900	35	246	3.33
(96,96,49)	300	35	246	1.11
(96,96,40)	300	60	309	1.11
(96,96,70)	300	10	177	0.74

Table 1: Simulation parameters for each grid level

Initialization and lateral boundary conditions

Initial and boundary conditions are obtained from the National Center for Atmospheric Research (NCAR) North American Regional Reanalysis (NARR) data. NARR is available at 32 km horizontally at 29 vertical levels (1000-100hPa; excluding surface) to force the WRF simulations at the coarsest grid. Lateral boundary condition forcing was applied at three-hour intervals and linearly interpolated in between. Relaxation towards the lateral boundary values was applied around the edge of the domain. Simulations were performed for 48 hours for the ramping events and for 168 hours for the week-long statistical analysis. Output was stored at ten-minute intervals and used to generate initial and boundary conditions for the nested grids.

Surface characteristics

WRF uses 33 land use categories (including ice and water) from USGS. The simulations used National Land Cover Database (NLCD) data at 30 m resolution which was mapped from the native 19 land use categories to the USGS 33 land use categories. The Noah Land Surface Model is used with initial moisture and soil temperature in four layers provided by NARR. Land-atmosphere coupling effects on the boundary layer, such as soil moisture initialization, can have a significant effect on simulation results, and at finer resolutions it may be necessary to use finer-scale soil initialization fields [1, 12, 11].

Turbulence and computational mixing

The standard Mellor-Yamada-Janjic scheme, which is appropriate for Reynolds-Averaged Navier-Stokes simulations, is used for the base case. Mellor-Yamada-Janjic (MYJ) is a one-dimensional prognostic turbulent kinetic energy scheme with local vertical mixing. A turbulent kinetic energy (TKE) closure of order 1.5 is used for the large-eddy-simulations performed. The TKE-1.5 model is useful in an LES context as long as the chosen length scale is proportional to the LES filter width, as it is in WRF. This closure is particularly useful for LES when a large fraction of the velocity scales are contained in the subfilter scales as with coarse resolution grids. The Dynamic Reconstruction Model (DRM) [1], which has been implemented into WRF [7], was also used as an LES closure in the simulations. The DRM is a “mixed model” which combines an eddy-viscosity component with a scale-similarity component and has shown improved representation of turbulence in prior work [1, 2].

COMPARISON WITH OBSERVATION DATA

Results comparing simulations with observation data for the two-day ramping events and the week-long runs are given in this section. The synoptically driven ramping event (Case S) is caused by a frontal passage just after significant precipitation and appears 10 hours into the simulation (3 hours after sunrise) at which point wind speed magnitude increases from 3 to 15 m/s over a 50 minute period. The locally driven ramping event (Case L) appears 20 hours into the simulation (7 hours after sunrise) at which point the wind speed magnitude increases from 2 to 14 m/s in just over 8 hours. The week-long simulation is performed for 7 days starting at the same time as the locally driven ramping event in the summer. The naming convention for runs is “(finest horizontal resolution)_(number of points in vertical)”, as shown in table 2. All the results are from one-way nesting and the Mellor-Yamada-Janjic turbulence model unless otherwise noted.

Run name	Configuration
300m_49	base case simulation, 4 one-way nested grids
300m_70	4 one-way nested grids with 70 vertical levels
300m_40	4 one-way nested grids with 40 vertical levels (chosen by WRF)
300m_49_2way	4 two-way nested grids
900m_49	3 one-way nested grids
2700m_49	2 one-way nested grids
8100m_49	single grid
300m_49_TKE1.5	4 one-way nested grids, with TKE1.5 closure on the finest grid
900m_49_TKE1.5	4 one-way nested grids, with TKE1.5 closure on the finest grid
100m_49_TKE1.5	5 one-way nested grids, with TKE1.5 closure on the finest grid
100m_49_DRM1	5 one-way nested grids, with DRM closure at reconst. level 1 on the finest grid

Table 2: Simulation configurations

Effects of resolution

Effects of horizontal resolution

The effect of horizontal resolution was investigated to determine if it is significant in simulating both Case S (synoptically-induced) and Case L (locally-driven) ramping events. We compare output from one-way nested runs to evaluate the impact of grid resolution. SODAR observations of the evolution of wind speed and wind direction 30 m above the surface during the ramping event are presented in figure 2 (left) for Case S and at 90 m in Figure 2 (right) for Case L (heights chosen are different due to observational data gaps). The observations are compared to one-way nested run resolutions: 300m_49 (4 grid nests), 900m_49 (3 grid nests), 2700m_49 (2 grid nests), and 8100m_49 (1 grid). Regardless of resolution, the simulation for Case S shows an increase in wind speed associated with this ramp that is slightly more gradual than was observed by the SODAR. The simulation for Case L also shows a more gradual increase in wind speed associated with the ramping event and does not capture the minimum wind speeds observed by the SODAR.

Figure 3 shows vertical profiles of the two ramping events. In general, the simulations better capture the wind speed after ramping rather than right before it. SODAR data are available at 10 m intervals from 20-200 m above ground level, however at elevations above 100m there is usually little data available. Figure 4 shows the vertical structure of the atmosphere in time for the simulations interpolated to the SODAR's vertical levels. Case L overpredicts wind speeds after the first 16 hours of the simulation and underpredicts before the 16 hours, while Case S seems to predict wind speeds more closely to the SODAR.

To further quantify the comparison, table 3 shows the root-mean-square errors (RMSE) and mean absolute errors (MAE) between the SODAR data and the simulations over 0-200 m above the ground, averaged over 48 hours. They are defined as:

$$MAE = \frac{1}{M} \sum_{j=1}^M \frac{1}{N} \sum_{i=1}^N |A_{i,j} - B_{i,j}|$$

$$RMSE = \sqrt{\frac{1}{M} \sum_{j=1}^M \frac{1}{N} \sum_{i=1}^N (A_{i,j} - B_{i,j})^2}$$

where M =number of time steps; N =number of vertical grid points; A =observation; and B =simulation data.

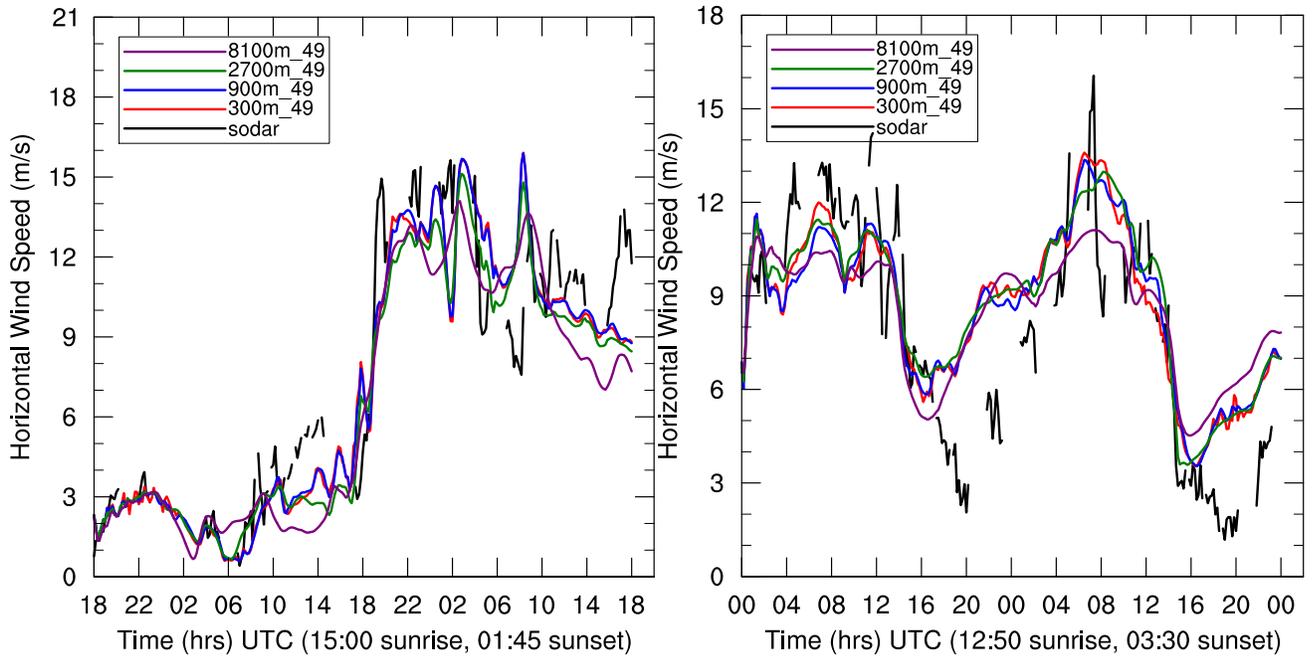


Figure 2: Total wind speed time series for Case S (left) and Case L (right) for different horizontal resolutions (grid nesting levels)

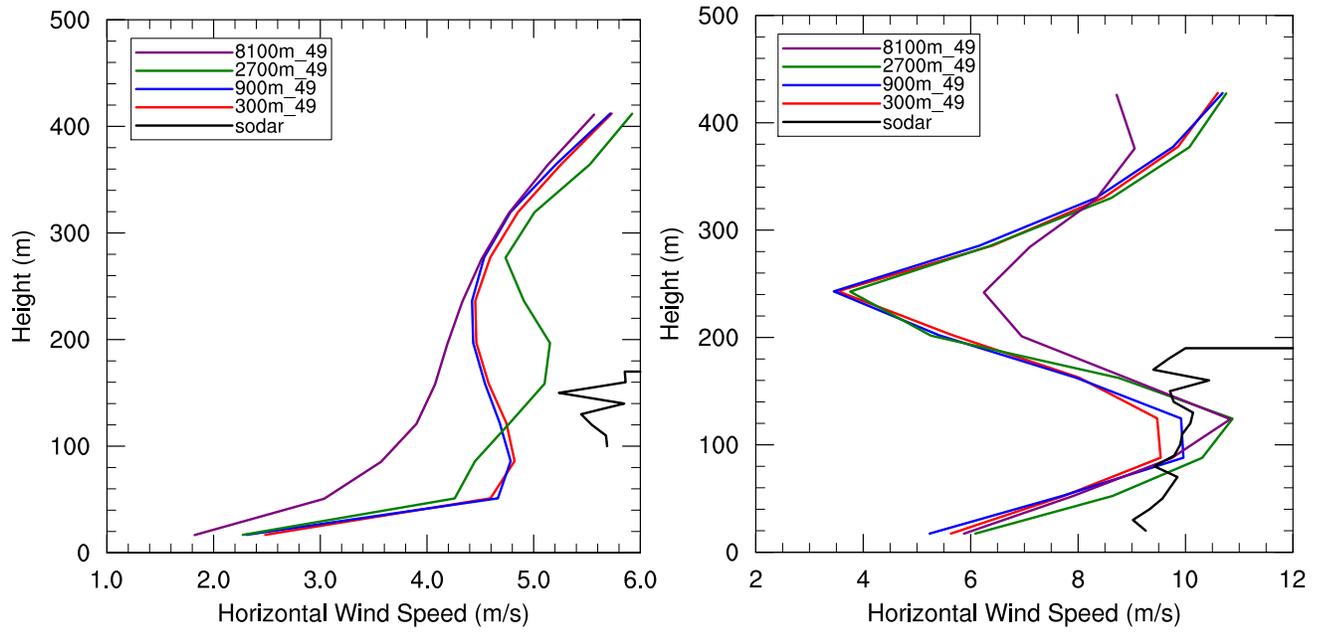


Figure 3: Total wind speed vertical profiles for Case S (left) and Case L (right) at a single time slice

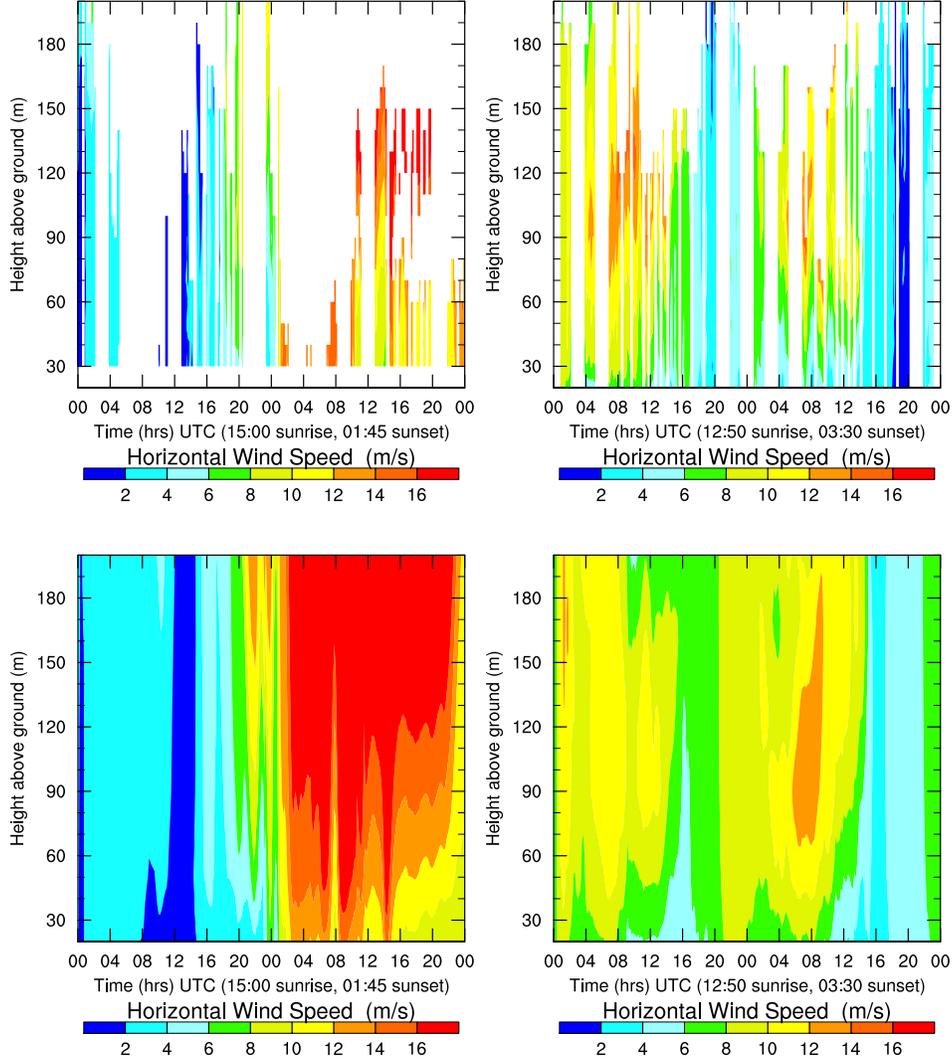


Figure 4: Vertical Structure of wind speed for Case S (left column) and Case L (right column) compared to SODAR (top row)

Resolution	Synoptically-Driven Case		Locally-Driven Case	
	MAE	RMSE	MAE	RMSE
300m_49	1.67	2.30	2.08	2.54
900m_49	1.62	2.23	2.05	2.48
2700m_49	1.63	2.16	2.01	2.47
8100m_49	1.77	2.30	2.23	2.69

Table 3: RSME and MAE for wind speed at different horizontal resolutions

The difference between the errors for the various grid resolutions is not very significant, as seen in the error metrics listed in table 3. The results are very similar across all resolutions with the finest resolution producing a slightly larger error than the two finer ones preceding it. While all the runs capture the increase

in wind speed during the ramping event itself, Case S appears to be much better captured and with slightly less error than Case L. Under these particular conditions relatively coarse horizontal resolution (2.7 km) appears to be ideal in capturing the shift in wind speed at the wind turbine hub heights ($\sim 80\text{-}100$ m) for both Cases S and L.

Effects of vertical resolution

Vertical spacing may have a significant effect on the accuracy of wind speed and direction predictions when wind shear is strong, e.g. in the case of nocturnal low level jets which can also be responsible for ramping events (see e.g. [10]). The vertical grids tested here using 4 one-way nesting levels were generated as follows: 300m_49 uses a *tanh* function to stretch the vertical grid spacing so that there is smaller spacing closer to the ground; 300m_70 uses specified grid levels that provide spacing of about 10 m in the lowest 200 m of the atmosphere and then stretched; 300m_40 was created by specifying 40 levels in WRF and using the default spacing (the minimum spacing was 60 m and average spacing was 309 m). Figure 5 shows the evolution of the wind speed and table 4 shows the RMSE and MAE between simulations and observations along with the average vertical spacing in the range of the observation data (20-200m). The grids all give similar results for this ramping event. The finely spaced vertical resolution is marginally better at capturing the ramping event than the coarse WRF defined levels. The *tanh* spaced 49 level simulation performs slightly better than the others in both ramping events, but even that is only by roughly 0.04 m/s. Once again the less computationally intense option appears adequate for predicting ramping.

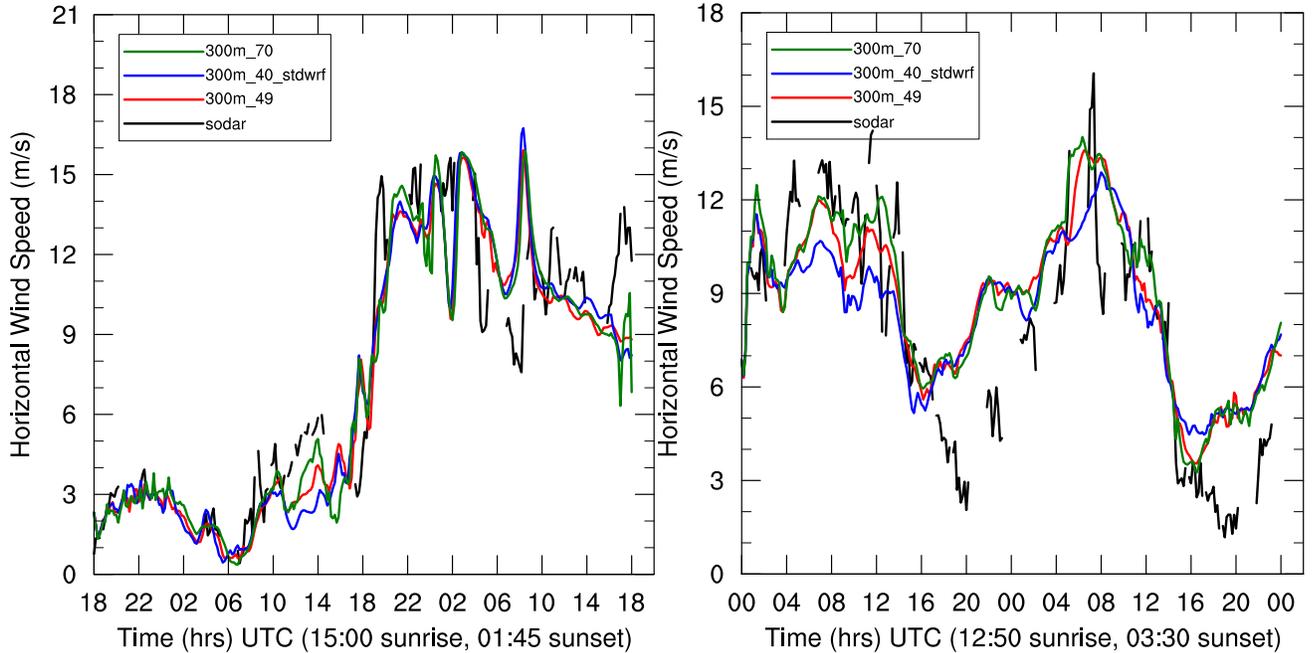


Figure 5: Total wind speed time series for Case S (left) and Case L (right) for different vertical resolutions

Resolution	Synoptically-Driven Case		Locally-Driven Case	
	MAE	RMSE	MAE	RMSE
300m_49	1.67	2.30	2.05	2.52
300m_40_stdwrf	1.73	2.39	2.07	2.59
300m_70	1.72	2.33	2.07	2.55

Table 4: RMSE and MAE for wind speed for different vertical resolutions

Difference between 1-way and 2-way nesting

With two-way nesting, the finer grid feeds back information to the coarser parent domain, thus updating the coarser grid with higher resolution fields and theoretically leading to a more accurate simulation result. In one-way nesting, the finer domains do not influence the parent domains at all. Two-way nesting is thought to be important when fine-scale features that are resolvable on the finer domain affect meteorological conditions on a larger scale. Under strong synoptic-scale forcing, it is likely that one-way nesting will be adequate for feeding information from the larger scales to the finer scales. This also seems to be the case here for the Case L (locally-driven forcing) as figure 6 shows, likely because of the simplicity of the topography. Figure 6 shows the time series of wind speed and direction for the 1-way nested base case, 300m_49, and the 2-way nested 300m_49_2way, each with the same four nesting levels. Table 5 shows the MAE and RMSE between simulations and observations. There is no significant difference between 1-way and 2-way nesting in WRF’s simulations compared to observations for both Cases S and L.

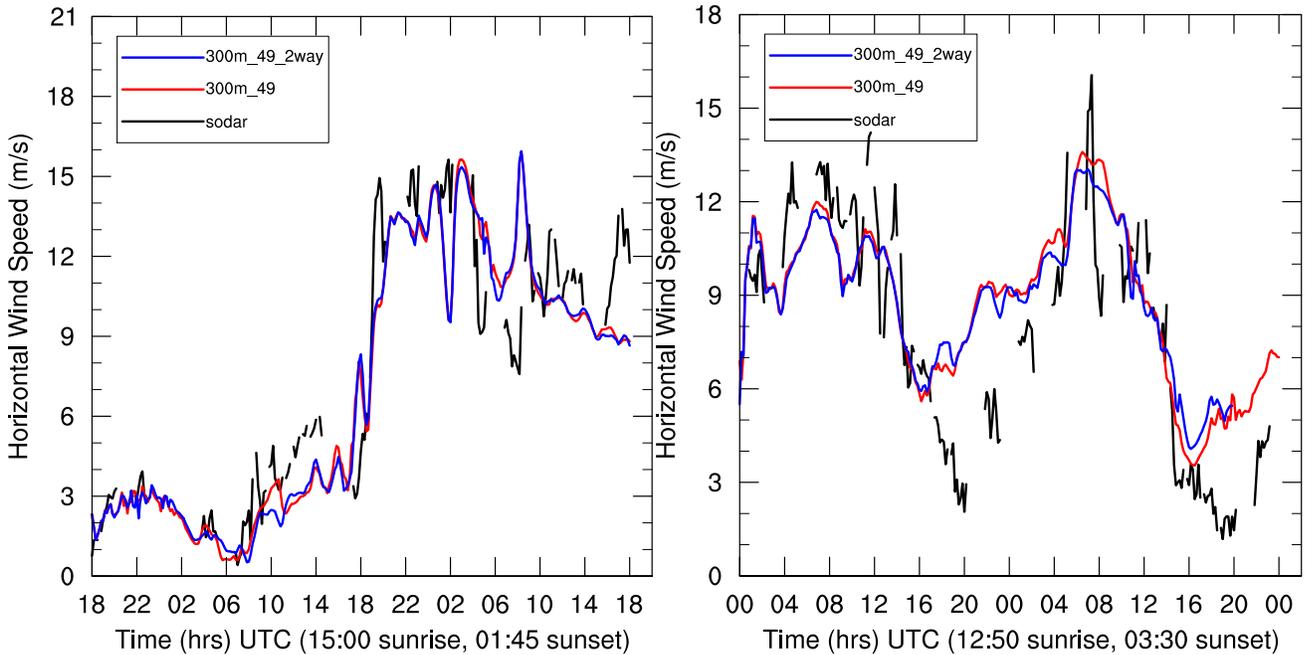


Figure 6: Total wind speed time series for Case S (left) and Case L (right) comparing nesting

Resolution	Synoptically-Driven Case		Locally-Driven Case	
	MAE	RMSE	MAE	RMSE
300m_49	1.67	2.30	2.08	2.54
300m_49_2way	1.68	2.31	2.07	2.52

Table 5: Root-mean-square errors (RMSE) and mean absolute errors (MAE) for wind speed comparing nesting for the two ramping events

Comparing RANS to LES closures

At horizontal resolutions finer than 1 km, it may be beneficial to transition to a large-eddy simulation closure within a nested simulation framework. WRF’s standard LES closure is the TKE 1.5 scheme. We performed simulations using the TKE 1.5 turbulence model for the finest domain level (with the coarser grids using Mellor-Yamada-Janjic) and compared to the Mellor-Yamada-Janjic model at all nesting levels. Both runs are 1-way nested because simulations with TKE 1.5 are unstable for 2-way nesting in this case (see discussion in

[8]). Figure 7 shows the TKE 1.5 model at different resolutions compared to the base case. Surprisingly our base case provides the best comparison to observation results. The TKE 1.5 closure tends to over predict the wind speeds in both ramping cases even at the finest resolution of 100 m. The results significantly improve for TKE 1.5 with finer resolution but are not better than the base case, as shown in table 6. Figure 8 shows how DRM level 1 compares to TKE 1.5 and the base case. The DRM closure does not show any improvement over the TKE 1.5 scheme at the same level. In WRF, the TKE 1.5 closure does not perform well with diffusion during the presence of strong heat fluxes for coarser grids as it uses an explicit scheme for calculating the diffusion. This is a possible reason that the TKE 1.5 closures do not perform as well at coarser resolutions. The over-prediction in wind speed decreases with finer resolution, until 100 m when it closely matches the estimates of the 300 m MYJ run. There does not seem to be great benefit in using LES at these fine scales because the finest grid very much relies on and is influenced by the RANS parent domains, which limit the fine grid's ability to vary very much from the parent.

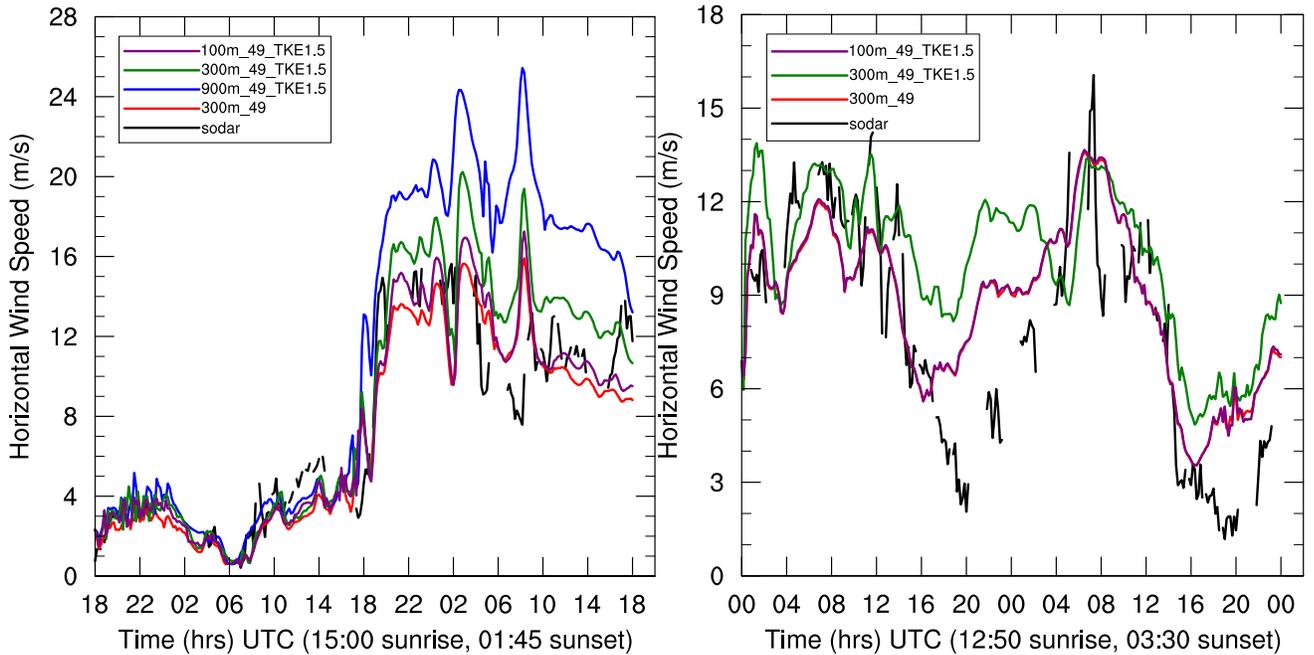


Figure 7: Total wind speed time series for Case S (left) and Case L (right) for TKE 1.5 closure at different horizontal resolutions

Resolution	Synoptically-Driven Case		Locally-Driven Case	
	MAE	RMSE	MAE	RMSE
300m_49	1.67	2.30	2.08	2.54
900m_49_TKE1.5	4.01	5.51	-	-
300m_49_TKE1.5	2.29	3.14	3.54	4.09
100m_49_TKE1.5	1.73	2.43	2.10	2.56
100m_49_DRM1	1.73	2.44	2.08	2.54

Table 6: Root-mean-square errors (RMSE) and mean absolute errors (MAE) for ramping events comparing turbulence closure

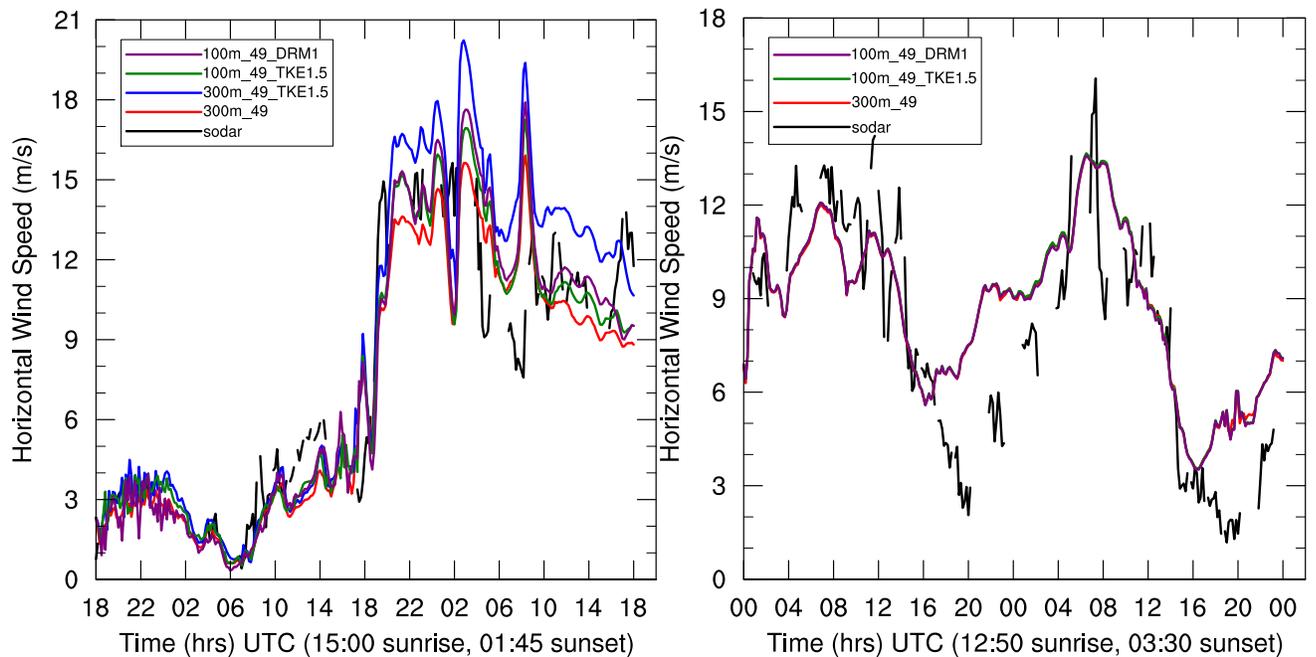


Figure 8: Total wind speed time series for Case S (left) and Case L (right) comparing DRM with TKE 1.5 closures

Comparing resolution with week-long simulations

A week-long simulation was performed for seven days starting on the same day as the locally-driven test case in the summer. The goal was to determine if grid resolution affects predicted disturbances of wind speed and turbulence intensity distributions for a specific site. Changes in long-term distributions could enable more detailed site evaluations for turbine placement.

Week-long results for wind speed

The results obtained from running week-long simulations for horizontal wind speeds fall in line with the shorter term 48 hour simulations. The finer resolution runs show little variation from the coarser resolutions as shown in the probability density function (PDF) in figure 9. The PDFs were created at 90 m height above ground and binned into 18 equally spaced bins that span the minimum and maximum range for each category at the SODAR location during a week of summer. Wind speeds were grouped into bins and each bin's probability was calculated so that all the bins added up to one hundred percent. They all show, along with the SODAR, that during this summer week, wind speeds are rather high and favorable for producing wind energy.

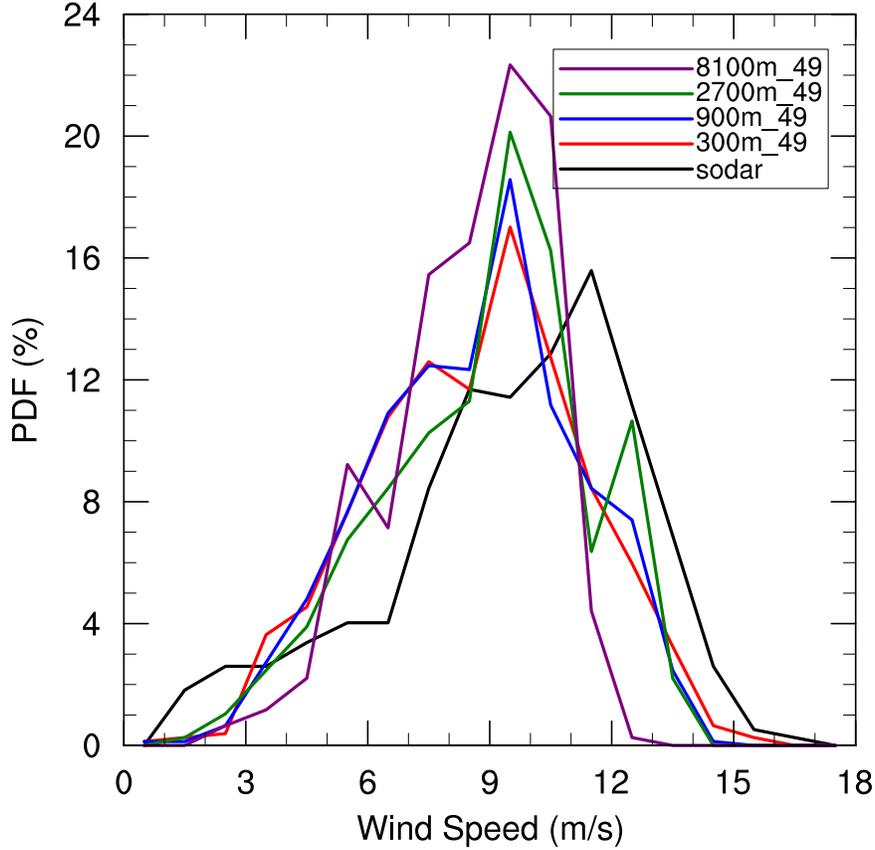


Figure 9: PDF for horizontal wind speed for week-long simulation

Week-long results for turbulence intensity

Turbulence intensity is important because if it gets too high (roughly greater than 20% of the mean wind speed) wind turbines will need to be switched off in order to avoid damage. Turbulence intensity is also vital to consider in determining potential wind farm locations. It is not adequate for a site to only have high wind speeds; it should have tolerable levels of turbulence intensity as well, otherwise the turbines will be offline for most of the day(s). The definition of turbulence intensity is:

$$TI = \frac{\sqrt{\sigma_u^2 + \sigma_v^2 + \sigma_w^2}}{U_{mean}}$$

where TI is turbulence intensity, U_{mean} is the average total speed over a time period, σ_u ; σ_v ; σ_w are the standard deviations of wind speed over that same time period in the East-West, North-South, and vertical directions respectively. The turbulence intensity for the SODAR was calculated by averaging the wind speed at every recorded second over a 10 minute period and calculating the standard deviation of each component of horizontal wind speed. The turbulence intensity for the simulations was calculated using WRF's output of TKE for MYJ:

$$\frac{\sqrt{TKE}}{U}$$

where TKE is the WRF output of turbulent kinetic energy for a specific time step and U is the total speed at the same time step. As figure 10 shows, there is not a large difference between the different resolutions as they match up to the SODAR. They all show similar results and estimate probability areas close to the

SODAR. The peaks for the turbulence intensity are slightly off, with the simulations predicting slightly lower turbulence intensity when compared to the SODAR. However, as the number of bins increases the peaks come closer together. This site seems to be rather ideal during this time of year for wind turbines, as turbulence intensity only exceeds 20% less than 4% of the week.

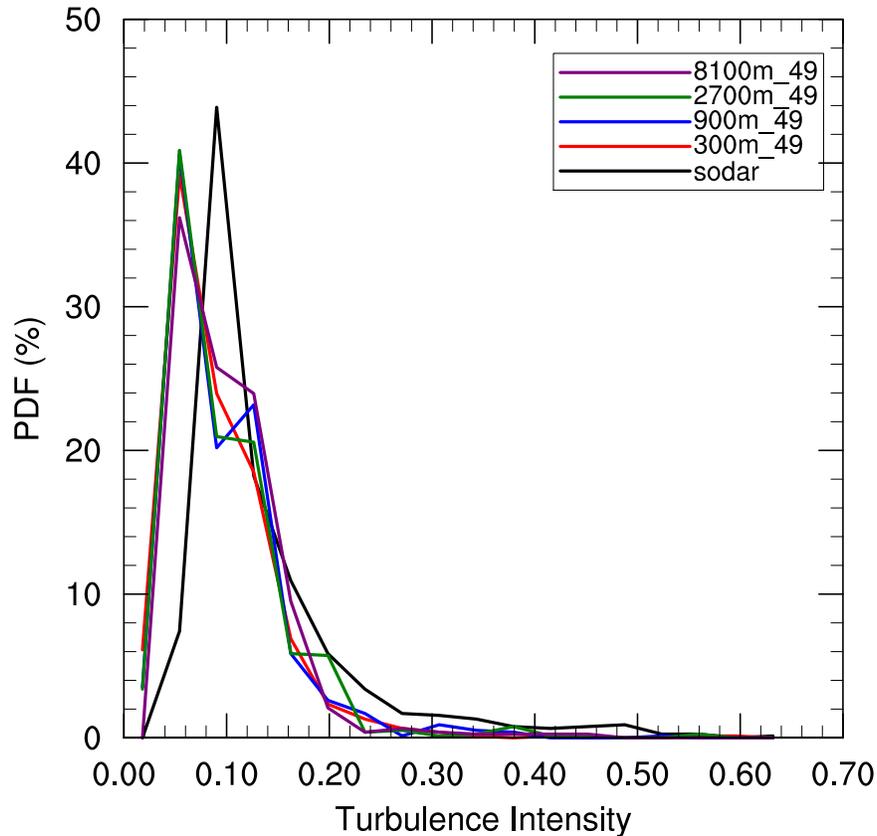


Figure 10: Turbulence intensity PDF for week long simulation

Topographic resolution analysis

A potential explanation for why the simulations show little change with higher resolutions is that the terrain of the wind farm is comprised only of a few rolling hills and flat farmland as shown in figure 1. Thus even the coarsest simulations used here can adequately capture the terrain. Figure 11 shows how all the resolutions, fine and coarse (except for the 8100 m case), capture the hill and general topography of the wind farm quite well. The coarsest domain (8.1 km) is the least effective at capturing the terrain and shows the greatest errors in the simulations discussed above.

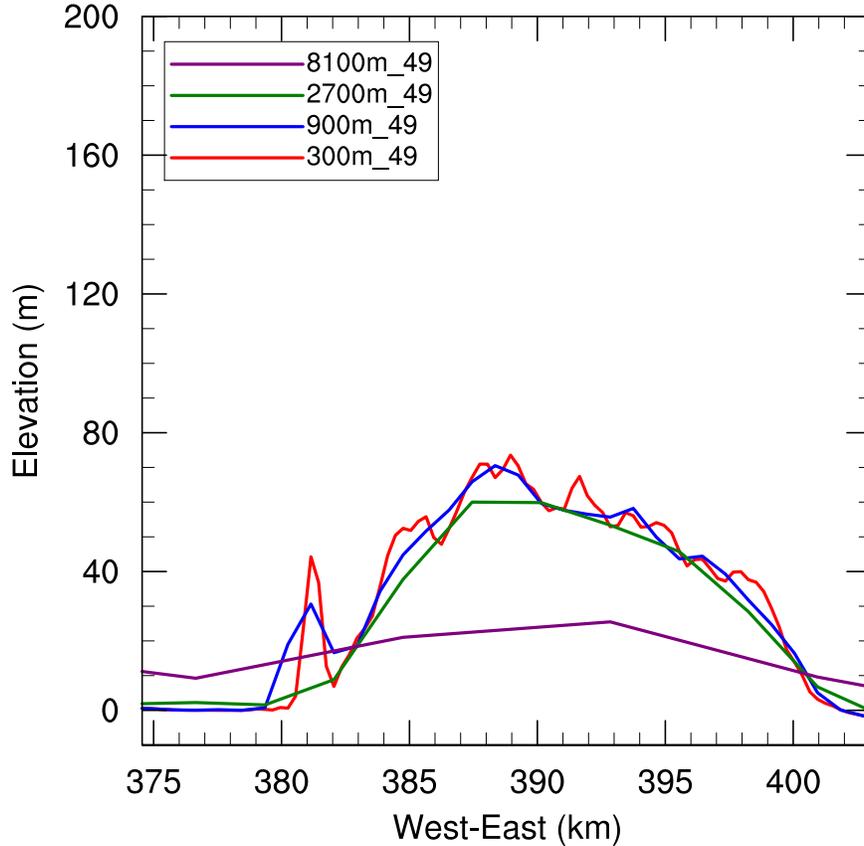


Figure 11: Terrain resolved by different horizontal resolutions

Land surface conditions and effects on results

Here we examine the effect of land surface forcing on the ramping event for Case L. Specifically sensitivity tests are designed to test the role of soil moisture in these wind transitions. In general, Case L is anticipated to be more sensitive to land-surface forcing. The Case L ramping event shows greater errors when compared to SODAR observations than Case S, especially in capturing the ramping at hub height, so there is certainly room for improvement in the simulations. Figure 12 (left) shows that when the soil is initialized as dry, 300m_49_2way_dry, it matches our base case fairly well by not capturing the minimum wind speeds but comparing well with the SODAR at higher wind speeds. The opposite is true for saturated soil, 300m_49_2way_sat. The simulation in that case captures the minima of the SODAR better while being out of phase with the maxima. The saturated case shows significantly reduced error when compared to the base and dry cases as shown in table 7. Figure 12 (middle and right) show that even a three degree difference in potential temperature or a 0.004 kg/kg difference in water vapor mixing ratio in the atmosphere can lead to drastic changes in wind speed as shown in figure 13. This suggests that replicating the proper land-surface conditions might be more important in capturing the ramping event in this case, rather than looking at differences in resolution or choice of turbulence closure. It is important to consider many other parameters in a simulation apart from resolution and turbulence closures such as land surface models and soil moisture initialization as these may yield greater reductions in error and better capture the ramping event.

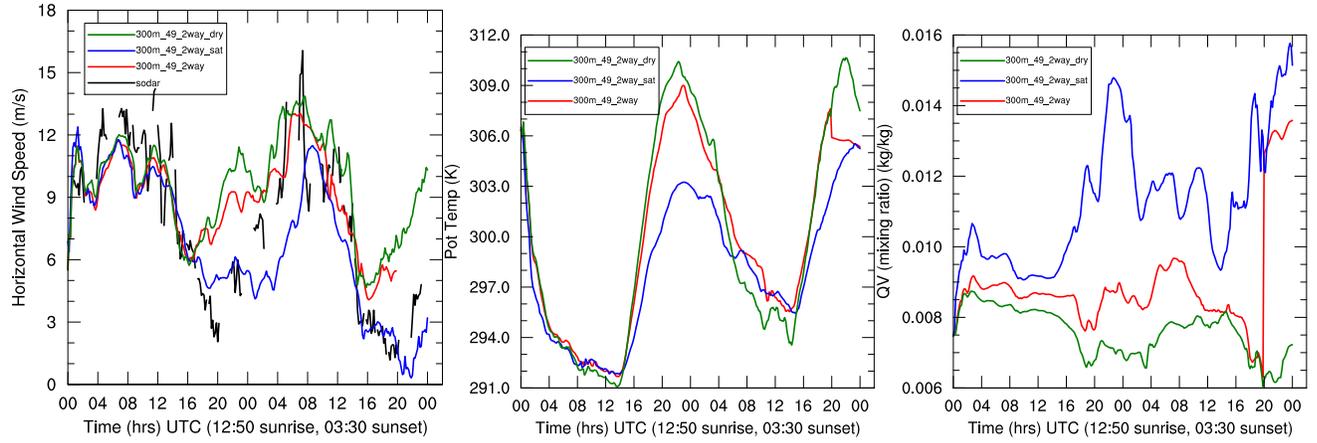


Figure 12: Land-surface variations for wind speed (left), potential temperature (middle) and water vapor mixing ratio (right) at 90 m above ground

Resolution	Locally-Driven Case	
	MAE	RMSE
300m_49_2way	2.07	2.52
300m_49_2way_sat	1.62	2.10
300m_49_2way_dry	2.66	3.24

Table 7: MAE and RMSE for Case L for different land-surface conditions

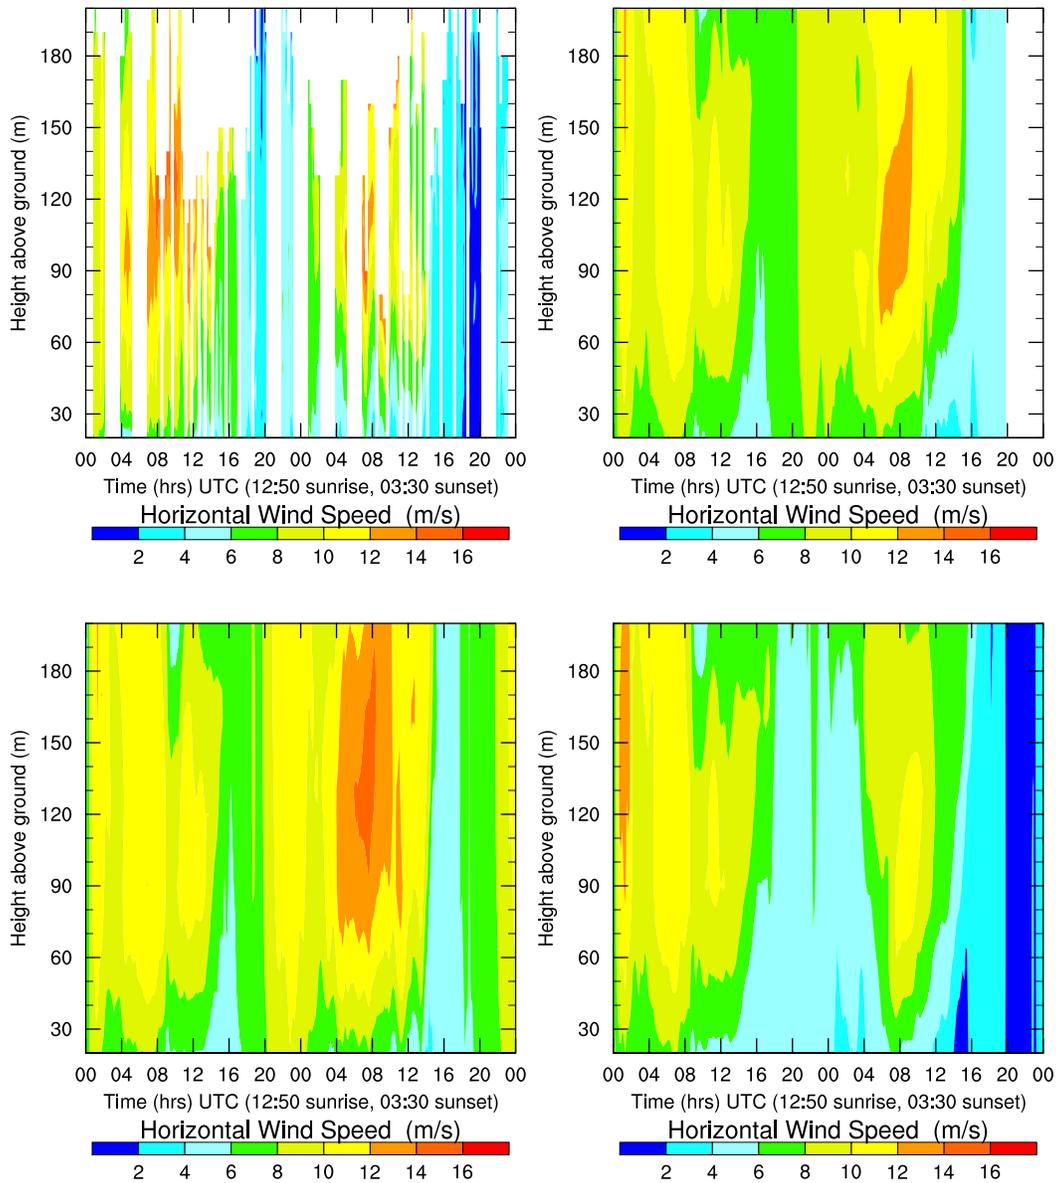


Figure 13: Wind speed time series contours for SODAR (top left), base case (top right), dry case (lower left) and saturated case (lower right)

The effects of external forcing

The source of initial and boundary condition data can significantly affect simulation results and should also be considered. Figure 14 shows how two different sets of external forcing data can lead to very different results. When the base case with NARR forcing was simulated for this day in spring, the ramping event was simulated eight hours too early. The simulation run with North American Model (NAM) forcing (300m_49_NAM) enabled the coarsest grid (8.1 km) to be eliminated because NAM forcing is at 12 km resolution as opposed to NARR forcing at 32 km. The new simulation captured the ramping event on time in this case, but not its magnitude. Table 8 shows that a 1 m/s improvement in error can be achieved by using different initialization and boundary conditions.

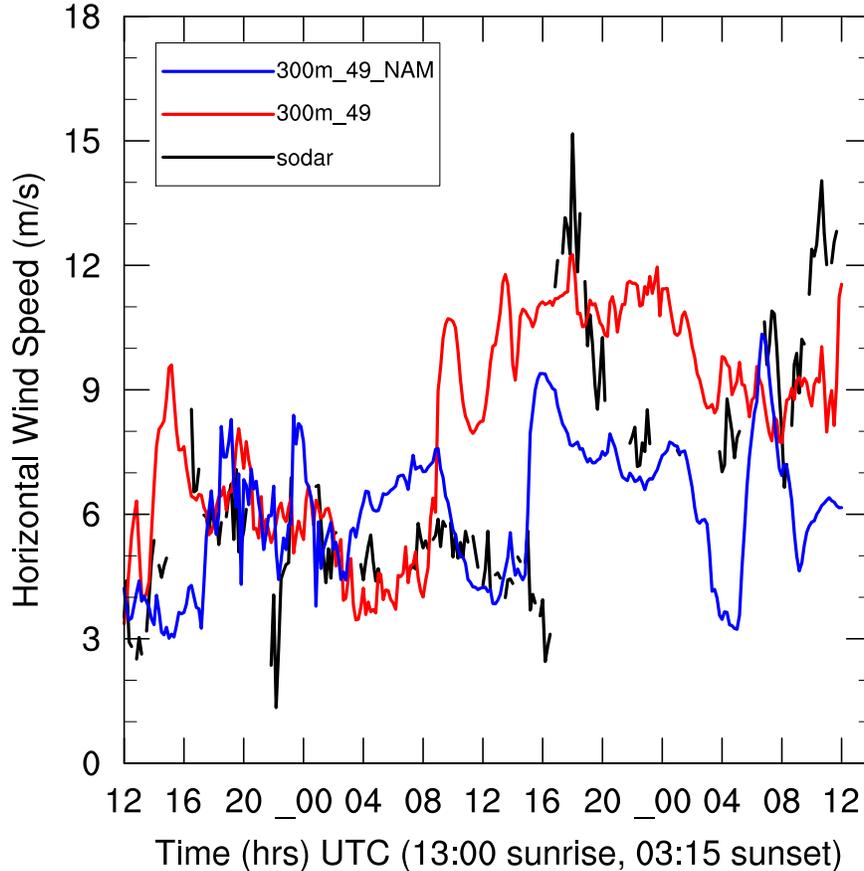


Figure 14: Horizontal wind speed time series comparing NAM and NARR forcing on spring day

Resolution	Locally-Driven Case	
	MAE	RMSE
300m_49	3.29	4.73
300m_49_NAM	2.38	3.06

Table 8: MAE and RMSE for comparing NAM and NARR forcing for a spring day

Conclusion

This paper presents simulation results using WRF to capture synoptically and locally-driven ramping events at a wind farm on the West Coast of the United States. The wind ramping was closely represented with all grid resolutions and grid nesting configurations tested. These results seem counter-intuitive based on general practical recommendations to use higher resolution, but given the simple terrain of this wind farm it is logical to conclude that resolution does not necessarily bestow the benefit of increased accuracy. Moreover, the LES turbulence closure does not show much benefit because it has difficulty at coarser resolutions with WRF. At very fine resolutions, the LES turbulence closures perform well, but there is little added benefit because at these resolutions the domain is too small to show results differing too greatly from RANS (which provides the lateral boundary forcing). It is of interest to note that the synoptically-driven event is better simulated than the locally-driven event; this is expected because the strong synoptic forcing usually overshadows other parameters. Land-surface parameters have more of an effect on the simulation results than resolution or

our options for turbulence closure. Initialization conditions could also have a significant effect on simulation results in some cases. Future work will be performed on more complex terrain to investigate the effects of resolution and more land-surface model parameters should be experimented with, such as soil moisture and landuse to determine just how sensitive any changes may be.

References

- [1] FK Chow, RL Street, M Xue, and JH Ferziger. Explicit filtering and reconstruction turbulence modeling for large-eddy simulation of neutral boundary layer flow. *JOURNAL OF THE ATMOSPHERIC SCIENCES*, 62(7, Part 1):2058–2077, JUL 2005.
- [2] Fotini Katopodes Chow and Robert L. Street. Evaluation of Turbulence Closure Models for Large-Eddy Simulation over Complex Terrain: Flow over Askervein Hill. *JOURNAL OF APPLIED METEOROLOGY AND CLIMATOLOGY*, 48(5):1050–1065, MAY 2009.
- [3] J Feedman, M Markus, and R Penc. Analysis of West Texas wind plant ramp-up and ramp-down events. *AWSTruewind report*, 2008.
- [4] GA Grell, S Emeis, WR Stockwell, T Schoenemeyer, R Forkel, J Michalakes, R Knoche, and W Seidl. Application of a multiscale, coupled MM5/chemistry model to the complex terrain of the VOTALP valley campaign. *ATMOSPHERIC ENVIRONMENT*, 34(9):1435–1453, 2000.
- [5] S Gronas and AD Sandvik. Numerical simulations of local winds over steep orography in the storm over north Norway on October 12, 1996. *JOURNAL OF GEOPHYSICAL RESEARCH-ATMOSPHERES*, 104(D8):9107–9120, APR 27 1999.
- [6] SR Hanna and RX Yang. Evaluations of mesoscale models’ simulations of near-surface winds, temperature gradients, and mixing depths. *JOURNAL OF APPLIED METEOROLOGY*, 40(6):1095–1104, 2001.
- [7] JD Mirocha, FK Chow, JK Lundquist, and KA Lundquist. Improved subfilter turbulence modeling for large eddy simulation using WRF. *7th Symposium on the Urban Environment, American Meteorological Society*, (13.1):4, 2007.
- [8] C.-H. Moeng, Jimy Dudhia, Joe Klemp, and Peter Sullivan. Examining two-way grid nesting for large eddy simulation of the PBL using the WRF model. *MONTHLY WEATHER REVIEW*, 135(6):2295–2311, JUN 2007.
- [9] William C. Skamarock and Joseph B. Klemp. A time-split nonhydrostatic atmospheric model for weather research and forecasting applications. *JOURNAL OF COMPUTATIONAL PHYSICS*, 227(7):3465–3485, MAR 20 2008.
- [10] Brandon Storm, Jimy Dudhia, Sukanta Basu, Andy Swift, and Ian Giammanco. Evaluation of the Weather Research and Forecasting Model on Forecasting Low-level Jets: Implications for Wind Energy. *WIND ENERGY*, 12(1):81–90, JAN 2009.
- [11] JL Williams, RM Maxwell, JK Lundquist, and S Wharton. On the use of a coupled variably-saturated groundwater flow - land surface model to initialize a fully coupled subsurface - land surface - atmospheric model for wind energy forecasting. *Eos Trans. AGU Fall Meeting Supplement*, 90(52, A31F-0183), 2009.
- [12] CS Woodward, RM Maxwell, JK Lundquist, J Mirocha, S Smith, and AF Tompson. Wind Energy Resource Assessment using Coupled Groundwater-Land-Surface Atmospheric Models. *Eos Trans. AGU Fall Meeting Supplement*, 90(52, A31F-0189), 2009.
- [13] JC Wyngaard. Toward numerical modeling in the “terra incognita”. *JOURNAL OF THE ATMOSPHERIC SCIENCES*, 61(14):1816–1826, JUL 2004.