



Task 1 Final Report (Revision 1)
for contract M11PC00017

Project Title:
**Prediction of Wind Energy Resources on the Outer Continental Shelf
with Weather Models**

Submitted to:
Bureau of Ocean Energy Management (BOEM)



Prepared by:
MMI Engineering & Atmospheric and Environmental Research

Revision 1 Submitted: November 8, 2011

(Original: October 19, 2011)

Table of Contents

| | |
|---|----|
| A. Executive Summary..... | 1 |
| B. Methodology | 3 |
| The Advanced Research WRF model..... | 3 |
| A 30-year reanalysis using NARR | 3 |
| The model domain | 4 |
| Choice of model physics..... | 8 |
| Running WRF | 11 |
| Choice of testing period | 11 |
| Validation | 14 |
| C. Test results | 17 |
| Radiation time-step | 17 |
| Damping layer..... | 19 |
| Boundary layer parameterization | 20 |
| D. Recommendation..... | 40 |
| Appendix | 42 |
| Location of vertical levels | 42 |
| Model time step | 43 |
| Choice of sampling interval..... | 43 |
| Boston simulations..... | 45 |
| Comparison to results from the North Sea..... | 47 |
| List of Acronyms and Abbreviations | 49 |
| References | 50 |

Table of Figures

| | |
|--|----|
| Figure 1: Map of the four WRF domains to be used in Task 2 with grid spacing of 500 m (blue), 1.5 km (green), 4.5 km (red), and 13.5 km (black). Also shown are the OCS area of interest (shaded), and the location of the two closest buoys (crosses)..... | 5 |
| Figure 2: Map of the four WRF domains with grid spacing of 500 m (blue), 1.5 km (green), 4.5 km (red), and 13.5 km (black). Note that this domain is shifted onshore to allow verification against surface observations in Phase 1..... | 6 |
| Figure 3: Map of buoy locations in the area of Delaware Bay from the National Data Buoy Center via Google Maps. | 7 |
| Figure 4: Wind rose for the OCS from the NARR historical reanalysis for the full calendar year 2010. Each slice is representative of a directional bin (labeled in degrees from north). The length of the slice is representative of the relative frequency (%), labeled along the x-axis. The colored portions correspond to a speed (m/s) bin for each direction (i.e., green is 2.5 – 4.2 m/s)..... | 12 |
| Figure 5: Wind rose for OCS from the NARR historical reanalysis for the dates used in our sensitivity testing: every eleventh day in 2010 starting on January 7. Each slice is representative of a directional bin (labeled in degrees from north). The length of the slice is representative of the relative frequency (%), labeled along the x-axis. The colored portions correspond to a speed (m/s) bin for each direction (i.e., green is 2.5 – 4.2 m/s)..... | 13 |
| Figure 6: Location of surface observations used to verify our sensitivity tests. MADIS observations are marked with circles. NWS stations are marked with triangles. Buoy 44009 is marked with the cross. The OCS area of interest is marked with the dashed lines and is the same as that in Figure 1. | 15 |
| Figure 7: Surface (10 m) wind speed from a 30-hour simulation initialized at 00 UTC on March 1, 2010 with 1 minute radiation updates. | 18 |
| Figure 8: Difference between surface (10 m) wind speed from 1 minute and 10 minute radiation updates from a 30-hour simulation initialized at 00 UTC on March 1, 2010. Note that the differences of up to 0.32 m/s correspond to roughly a 6% difference and are largely related to the positioning of wind gusts associated with turbulent mixing in the boundary layer..... | 19 |
| Figure 9: Difference between surface (10 m) wind speed from simulation with and without an upper damping layer from a 30-hour simulation initialized at 00 UTC on March 1, 2010. Note that the differences are largely related to the positioning of small-scale features and show no systematic bias..... | 20 |
| Figure 10: Model error (forecast – observation) versus the forecast wind speed (m/s) for all forecasts using the MYJ boundary layer parameterization. The blue line is the least-squares trend indicating that the wind speed forecasts appear to be roughly double the observations. | 21 |
| Figure 11: Model error (forecast – observation) versus the forecast wind speed (m/s) for all forecasts using the MYNN boundary layer parameterization. The blue line is the least-squares trend indicating that the wind speed forecasts appear to be roughly double the observations. | 22 |

| | |
|--|----|
| Figure 12: The mean wind speed error (forecast – MADIS observations) for each day 11 th day of the calendar year 2010 (i.e. those modeled) with the MYJ boundary layer parameterization..... | 24 |
| Figure 13: The mean wind speed error (forecast – MADIS observations) for each day 11 th day of the calendar year 2010 (i.e. those modeled) with the MYNN boundary layer parameterization..... | 25 |
| Figure 14: The average forecast, error, and observed wind speed by hour of the day averaged across all simulations with the MYJ boundary layer parameterization. | 26 |
| Figure 15: The average forecast, error, and observed wind speed by hour of the day averaged across all simulations with the MYNN boundary layer parameterization..... | 27 |
| Figure 16: Histogram of model forecasts using the MYJ boundary layer parameterization from the MADIS observation locations. The histogram was built using a bootstrapping method sample equally from all hours of the days to remove a sample bias inherent in 30-hour forecasts..... | 28 |
| Figure 17: Histogram of model forecasts using the MYNN boundary layer parameterization from the MADIS observation locations. The histogram was built using a bootstrapping method sample equally from all hours of the days to remove a sample bias inherent in 30-hour forecasts..... | 29 |
| Figure 18: Histogram of MADIS observations taken from the MYJ match pair records. The distribution from the MYNN match pair records is nearly identical as the observations are drawn from the same data so only one distribution is presented. The histogram was built using a bootstrapping method sample equally from all hours of the days to remove a sample bias inherent in 30-hour forecasts. Note the clear maximum at the lowest wind speeds. | 30 |
| Figure 19: Histogram of model forecast errors from the MYJ simulations. The histogram was built using a bootstrapping method sample equally from all hours of the days to remove a sample bias inherent in 30-hour forecasts. Note the maximum is shifted to a 2-3 m/s bias..... | 31 |
| Figure 20: Histogram of model forecast errors from the MYNN simulations. The histogram was built using a bootstrapping method sample equally from all hours of the days to remove a sample bias inherent in 30-hour forecasts. Note the maximum is shifted to a 2-3 m/s bias..... | 32 |
| Figure 21: Histogram of all surface wind speed observations at buoy 44009 in calendar year 2010..... | 33 |
| Figure 22: Histogram of wind speed error (forecast – observation) at buoy 44009 for simulation using the MYJ boundary layer parameterization. | 34 |
| Figure 23: Histogram of wind speed error (forecast – observation) at buoy 44009 for simulation using the MYNN boundary layer parameterization. | 35 |
| Figure 24: Histogram of all surface wind speed observations at Georgetown, DE (KGED) in calendar year 2010..... | 36 |
| Figure 25: Histogram of all surface wind speed observations at Wildwood, NJ (KWWD) in calendar year 2010. | 37 |
| Figure 26: Photograph of the Georgetown, DE (KGED) ASOS station. Courtesy of the National Climatic Data Center. | 38 |

| | |
|--|----|
| Figure 27: The DEOS station in Ellendale, DE. Photo from the 17 January 2008 DEOS newsletter available at http://www.deos.udel.edu/news/011708.html | 39 |
| Figure 28: A photo of a 3-meter discus buoy from NOAA courtesy of the National Data Buoy Center. The buoy shown is (probably) not buoy 44009, but rather an identically equipped unidentified buoy. The original image is available at http://www.ndbc.noaa.gov/images/stations/3m_mini.jpg | 39 |
| Figure 29: Distribution of vertical levels. | 43 |
| Figure 30: Wind rose for OCS from the NARR historical reanalysis for the dates used in our sensitivity testing: every seventh day in 2010 starting on January 7. Each slice is representative of a directional bin (labeled in degrees from north). The length of the slice is representative of the relative frequency (%), labeled along the x-axis. The colored portions correspond to a speed (m/s) bin for each direction labeled in the upper-left (i.e., green is 2.5 – 4.2 m/s). | 44 |
| Figure 31: Wind rose for OCS from the NARR historical reanalysis for the dates used in our sensitivity testing: every fourth day in 2010 starting on January 7. Each slice is representative of a directional bin (labeled in degrees from north). The length of the slice is representative of the relative frequency (%), labeled along the x-axis. The colored portions correspond to a speed (m/s) bin for each direction labeled in the upper-left (i.e., green is 2.5 – 4.2 m/s). | 45 |
| Figure 32: Locations of MADIS observations used to verify the Boston simulation.. | 46 |
| Figure 33: Domain-averaged wind speed forecasts (red) and observations (blue) from the MADIS network (left) and official NWS stations (right) for a 10-day period in August 2010 from the Boston simulations. | 47 |
| Figure 34: Wind speed histogram from the Fino 1 platform (Beeken et al., 2008).. | 48 |

Table of Tables

| | |
|---|----|
| Table 1: Summary of nested domains. The colors given in the first column correspond to the model domains shown in Figure 1 and Figure 2..... | 8 |
| Table 2: Root mean square error $RMSE = (\text{mean}((\text{forecast} - \text{observed})^2))^{1/2}$, and mean error -- bias = $\text{mean}(\text{forecast} - \text{observed})$, for both boundary layer parameterizations verified against MADIS observations, buoy 44009, and official surface stations at Georgetown, DE (KGED), Wildwood, NJ (KWWD), and Ocean City, MD (KOXB). The values given for buoy 44009 included the actual observations at 5 m height and those "corrected" to 10 m height in parentheses. | 23 |
| Table 3: Vertical η levels used in the WRF simulations..... | 42 |

A. Executive Summary

As the government agency responsible for issuing offshore leases, the Bureau of Ocean Energy Management (BOEM)¹ has a duty to understand the wind energy resource potential of specific geographic regions. Offshore on the Outer Continental Shelf (OCS), where observations are typically extremely sparse, numerical weather modeling provides a method to gain insight into typical wind patterns to quantify wind energy potential. A wind resource climatology (herein defined using a standard 30-year window) will be useful in appraising a potential long-term lease on the OCS offshore of Delaware Bay.

The goals of this study are divided into two separable Tasks:

1. Test the numerical weather model to simulate historical weather conditions, validate the test results, and rank the tested model configurations
2. Predict long-term wind energy resources for an area on the OCS, and report the results in a format that can be applied to calculate probabilistic electricity generation estimates.

The final goal of the study is to use numerical weather models to produce a 30-year long climatology with 10-minute temporal and 500 m spatial resolution for an area on the OCS. Task 1 is essentially a sensitivity study using historical observations to test and recommend the optimal model configuration from a number of available choices. In Task 2, the recommended model configuration will be applied to develop the full 30-year time series. The report summarizes the results of Task 1 and presents recommendations for Task 2. The report is presented for the reference of a technical audience, but an effort is made to present brief explanations to make the report as accessible to a general audience as is practicably possible.

In Task 1 we have run a number of numerical weather simulations to represent typical conditions using the Weather Research and Forecasting (WRF) model (Skamarock and Klemp, 2008). The WRF model was chosen based on its broad choices of model physics and widespread community use. A model configuration that reproduces the wind speed climatology with little bias on the OCS (compared to measurement from the one available offshore buoy) is presented in this report. Most choices of model physics options are found to have little effect from a wind resources perspective. The recommended model configuration shows a bias that is smaller than the margin of error of the offshore observations used for verification. Accordingly the simulations were statistically the same as the actual surface observations.

While this configuration performs well within the scope of the project on the OCS, further study is needed for any future studies seeking to develop an onshore climatology. Significant data quality issues were identified for non-standard land-based observations collected from a network of unofficial weather stations. Such

¹ Previously, BOEMRE, the Bureau of Ocean Energy Management, Regulation, and Enforcement.

issues are beyond the scope of this study and do not affect the results of the study, but are presented as a caution for future studies that may seek to apply this methodology in other regions.

In Task 2, we will take the recommended model configuration and use it to develop a full 30-year climatology of wind resources on the OCS offshore of Delaware Bay. The climatology will be built by matching a historical reanalysis dataset, which will be explained later, with the most similar modeled time. This complete climatology can serve as a guide for the feasibility of an economical wind energy development on the OCS.

In this report we summarize the results of Task 1. In Section B, we discuss the methodology used in formulating and running our WRF simulations. Our results are validated against surface observations in Section C. The model configuration that is recommended to construct the full climatology is presented in Section D. The Appendix provides detailed technical discussions on select topics for the interested reader. A list of acronyms is given after the Appendix.

B. Methodology

The Advanced Research WRF model

The WRF model is an open-source (i.e., free to use and modify) weather model developed with contributions from over 150 organizations and universities. WRF is by far the most widely used and supported research and forecasting weather model in the world with a community of over 6,000 registered users. WRF is used for research purposes by a number of universities and government organizations. Both civilian and military government agencies use WRF for operational forecasting. In addition to forecasting, specialized versions of WRF exist for air pollution dispersion modeling, atmospheric chemistry, and wild fire prediction, etc. For the purposes of predicting wind resources, we must identify a version of the model to use, select a number of model configurations and dates to be tested, and verify the model forecasts against observations.

The WRF model has two separate cores that the user must choose between: The Nonhydrostatic Mesoscale Model (NMM) and the Advanced Research WRF (ARW). Each core evolved from a different predecessor model and uses different model grids and techniques to solve the equations that govern the atmosphere. Whether the ARW or NMM produces more accurate forecasts is seemingly a matter of the user's perspective and is a question that we will avoid.

The NMM was developed by and is used operationally at the National Centers for Environmental Prediction (NCEP, part of the National Weather Service - NWS) with a focus on speed of execution, numerical stability, and forecast consistency (Janjic, 2003). The NMM is tuned to run much faster than the ARW and is used primarily for short-term (3 days or less) forecasting. Specialized variants of the NMM are used for specific forecasting applications such as the Hurricane WRF.

The ARW, the version used for all simulations in this study, was developed by the National Center for Atmospheric Research (NCAR) and is used much more widely in research applications. However, it is still used operationally by groups including the Air Force Weather Agency (AFWA). The ARW supports much more numerous and complex numerical and physics options and is much better suited for high-resolution modeling. We are using the ARW largely for practical reasons related to its extensive list of supported physics, level of customizability, and ease of use and post-processing. Technical details on the ARW can be found in Skamarock and Klemp (2008).

A 30-year reanalysis using NARR

To construct a time series spanning 30 years, we must use a consistent dataset to initialize our simulations and build a full climatology. A reanalysis is a climate data set that is developed by combining weather modeling with observations to create a best-guess estimate of historical weather patterns. In that sense, weather model predictions are automatically 'reanalyzed' using actual observations to correct the model forecast. While it is tempting to simply simulate

the entire 30-year history using WRF, it is not necessary to simulate such a long period to determine the optimal model configuration. Instead the tested dates are drawn from the reanalysis to represent typical conditions.

The reanalysis most applicable to our purposes is the North American Regional Reanalysis (NARR) from NCEP (Mesinger et al., 2006). The NARR data is available on Lambert conformal projection (which approximates the Earth's surface as a cone), with a grid spacing of approximately 32 km in both latitude and longitude, and on 29 pressure levels in the vertical. NARR data is available back to 1979 in 3 hour increments and is actively being updated to include new data. The NARR dataset was used to prescribe the initial and boundary conditions for our numerical weather modeling.

For comparison, the alternative NCEP Climate Forecast System Reanalysis (CFSR) is available at a nominal 38 km resolution (Saha et al., 2010). The main advantage of the CFSR is its global coverage compared to the NARR, which is limited to North America.

The number of available physics packages that meet the customer needs led to the choice of a select number of tested model configurations. With this selected set of model configurations, the total number of days tested with each configuration was determined based on the total available computational resources per configuration divided by the number of computer hours per test. For example, given an arbitrary allocation of 100 units of computer time and two configurations to test, we have 50 units of computer time per configuration. If each day tested takes 2 units of computer time, then we can test 25 days for each configuration.

The model domain

Our study of the OCS requires forecasts with a 500 m grid resolution in the horizontal and 30 m in the vertical. In order to downscale the coarse NARR data, multiple nests are necessary to propagate the outer boundary conditions from a horizontal scale of approximately 32 km down to the 500 m fine-scale model grid while still maintaining physical balance. In this sense, a nest is a model grid that is placed within a coarser model grid for the purpose of producing higher resolution forecasts using information from the coarser grid. Wide boundaries between the edges of inner and outer nests are necessary to isolate the inner nests from artifacts of the outer boundary. This gives the atmospheric flow passing from the outer boundary toward an inner nest time to equilibrate within a domain before reaching the boundary of the inner nest.

To produce our fine-scale data we ran a triply nested WRF simulation (one with 4 total domains) to downscale the NARR data to the desired fine-scale grid. Our inner nest has a horizontal resolution of 500 m and contains 156 points in longitude and 196 points in latitude. A rectangular innermost domain stretching from the coast of the Delmarva Peninsula and southern New Jersey is necessary to capture the area of the OCS of interest for this study. The outer domains are then added to expand beyond this inner domain.

A 3:1 grid ratio between inner and outer nests is well established to be the optimal ratio for nesting WRF domains. This allows every third point on the inner nest boundary to fall on top of a point on the outer nest with two inner points in

between every outer point. This configuration was required in WRF's immediate predecessor for two-way nesting and so it remains the most widely studied and used setup. (A two-way nest is one where the inner and outer nests pass information back and forth to each other. In contrast, in one-way nesting the outer nest passes information to the inner nest, but the outer nest is not influenced by the inner nest.) Outside of the innermost nest, three additional domains with grid spacing of 1.5 km, 4.5 km, and 13.5 km telescope out with each domain adding 50% of the inner domain's length on each side. The domain that will be used in Task 2 to construct the full 30-year climatology is shown in Figure 1.




Figure 1: Map of the four WRF domains to be used in Task 2 with grid spacing of 500 m (blue), 1.5 km (green), 4.5 km (red), and 13.5 km (black). Also shown are the OCS area of interest (shaded), and the location of the two closest buoys (crosses)

In order to capture additional land-based surface observations, the innermost model domain was shifted onshore for use in the sensitivity study. The slightly shifted inner domain used in the sensitivity study of Task 1 is shown in Figure 2.




Figure 2: Map of the four WRF domains with grid spacing of 500 m (blue), 1.5 km (green), 4.5 km (red), and 13.5 km (black). Note that this domain is shifted onshore to allow verification against surface observations in Phase 1.

Note that only two instrumented buoys for measuring offshore wind (Figure 1, marked by the crosses) are present on the OCS at any point in the 30-year climatology. Both buoys are outside of the boundaries of the “area of interest” but are included nearby in the innermost model nest. As we would like to have more data for verification purposes in Task 1, we chose to shift the inner model domain toward the west (Figure 2). This brought the inner nest slightly onshore and allowed for a more robust statistical verification. As we will be using the same initial and boundary conditions, we do not expect that the shift will have any significant impacts on the validity of the forecasts as the initial data and physics will be unchanged.

The downscaled runs to be used in Task 2 will be computed using the domains shown in Figure 2. Ideally we would have liked to capture as much

verification data as possible. Additional shifting of the inner domain to include additional buoys is not possible as buoy observations in the region are very sparse and what buoys do exist are almost exclusively near-shore (Figure 3).

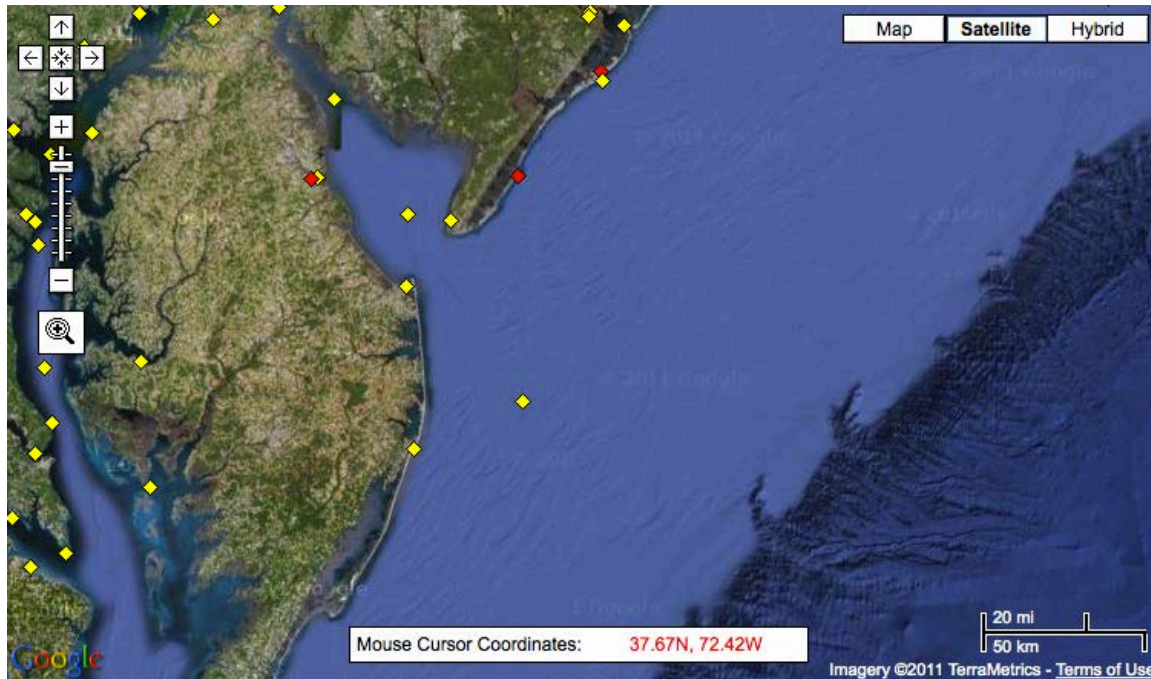


Figure 3: Map of buoy locations in the area of Delaware Bay from the National Data Buoy Center via Google Maps.

Whenever performing numeric simulations, it is important to consider the limits of numerical models when interpreting results. Not all physical processes are fully resolved as the resolution of the domain is much coarser than that of the finest physical processes. We expect our simulations to capture the large-scale sustained winds, which are defined by the World Meteorological Organization as the 10-minute average. We do not expect to fully capture details of processes like wind gusts with time scales of less than 5 minutes. We do not consider this a significant limitation, as the final goal of the study is a time series with 10-minute temporal resolution.

When considering WRF simulations, it is important to remember that physical processes of a scale less than seven times the horizontal model resolution are not well represented by the model (Skamarock, 2004). The implication for our application is that processes with characteristic scales shorter than 3 km to 4 km are not well resolved. This means that we should be able to capture most dynamics nearly down to the scales associated with individual thunderstorm cells. This should easily capture the dynamics associated with large scale wind patterns, but will miss some of the fine scale details associated with gustiness. If we assume a representative 10 m/s wind and the smallest resolve scale of 3 km, then wind variations on time scales of less than 5 minutes will not be fully resolved.

We have focused our simulation on the altitudes where wind generation occurs in the first few hundred of meters above the ocean's surface. This is the well-

mixed region of the atmosphere nearest the surface, called the boundary layer. WRF allows the user the option to specify the placement of vertical levels in the model. In our simulations we use a vertical resolution of approximately 30 m in the boundary layer, the same order as the required sampling levels (30, 60, 90, 120, 210 meters above the surface), and space the model levels further apart aloft. We use a total of 43 vertical levels as explained in Appendix: Location of vertical levels. As model coordinates are actually based on the mass of air in the column above the surface, and not on altitude, the best we can do is try to get the model levels to match the desired vertical levels in a “typical day”.

As the large-scale conditions are being forced by the NARR at the boundaries, we can significantly coarsen the vertical resolution in the middle and upper troposphere², where vertical gradients become weaker. We do not need high vertical resolution in our model to capture the details of the atmospheric flow aloft; the middle and upper troposphere are beyond our region of interest and will largely be driven by external forcing and not the model physics itself on the scale with which we are concerned.

Because we are focusing our study in detail in the lowest level of the atmosphere where turbulence is most prevalent, we need to ensure that our model is taking small enough steps in time so that it is stable. The fine vertical resolution in the boundary layer necessitates a lower time-step than would typically be used for a numerical simulation of this horizontal resolution. The time-step controls how far in time we progress the simulation every time we solve the governing equations. A brief discussion of choosing a model time step is given in the Appendix: Model time step. A summary of the domain sizes and time steps used is shown in Table 1.

Table 1: Summary of nested domains. The colors given in the first column correspond to the model domains shown in Figure 1 and Figure 2.

| <i>Grid spacing (km)</i> | <i>Number of points West-East</i> | <i>Number of points South-North</i> | <i>Time step (seconds)</i> |
|-------------------------------------|--|--|---------------------------------------|
| 13.5 - black | 49 | 61 | 45 |
| 4.5 - red | 73 | 91 | 15 |
| 1.5 - green | 103 | 136 | 5 |
| 0.5 - blue | 157 | 193 | 1.67 |

Choice of model physics

Choosing the proper model physics options from all of the possible combinations can be an overwhelming challenge. Having in-depth experience with weather models is essential to know which model settings should be used in which situations. In our case this challenge is aided by the specific requirements of the project. Not all options are recommended for real data cases. Not all packages provide all of the needed information. Some packages are so computationally

² The troposphere is the lowest part of the earth’s atmosphere (up to 8 – 12km), in which all processes relevant to weather take place.

intensive that they are not practical or necessary given our needs. Some options are designed to account for features, like urban land-use, which are not present in the area that we are modeling.

WRF simulations require the user to specify the following physics packages and parameterizations:

- A. Cloud microphysics: This controls how water in its various phases (rain, vapor, suspended droplets, ice, etc) are treated
- B. Radiation: This controls the energy balance in the model. Specific packages are required for incoming and outgoing radiation, plus the user must specify how often to call updates to these parameterizations.
- C. Surface layer: this specifies how the surface (ground and ocean) interacts with the atmosphere at its boundary
- D. Land surface: this specifies how the land is treated. In large part this selects what land-use (i.e., forest, grassland, urban, etc) classification system to use.
- E. Boundary layer: this parameterization propagates the results of the surface layer scheme to the rest of the atmosphere.
- F. Cumulus option: this specifies how to handle cumulus clouds that are too small to resolve in the model.
- G. Mixing options: there are multiple flags to specify how to control the mixing of heat, moisture, and kinetic energy between grid cells to account for the effect of unresolved turbulence
- H. Damping option: this specifies how to handle the upper boundary of the model.

Numerical diffusion is a way to account for fine-scale processes such as turbulence, molecular diffusion, and viscous friction that occur on scales that are too fine to be resolved in the model. In practice this is applied by averaging (i.e. smoothing) values across adjacent model points. Without diffusion's removal of small-scale features, the model would accumulate kinetic energy on its smallest resolvable scale and become very "noisy." As our simulations are initialized with real data and will need a parameterization for the boundary layer rather than simulating it directly, numerical diffusion is best performed along model surfaces. The horizontal diffusion coefficient (i.e. a measure of how much to smooth) is calculated by a two-dimensional Smagorinsky scheme based on the degree of horizontal deformation (Smagorinsky, 1963). The vertical diffusion (smoothing in the vertical) is handled by the boundary layer scheme.

The damping option at the top of the model domain specifies how energy reaching the model's upper boundary is absorbed or reflected. It essentially controls whether or not a sponge will absorb any kinetic energy reaching the top of the model. Without a damping layer, the model top is treated as a "rigid-lid" where up-going kinetic energy is reflected downward. It is very unlikely that the upper boundary will have noticeable effects near the surface where we are focused. At the altitudes of concern for this study, the effects of the damping were found to be negligible.

The radiation scheme in weather models handles the energy balance between heating from sunlight and cooling by infrared radiation into space. It also takes into account the role of clouds reflecting and scattering energy in both

directions. We use the Rapid Radiative Transfer Model (RRTMG) for both the longwave (outgoing) and shortwave (incoming) schemes. The RRTMG package is developed and maintained at AER and is distributed publically with WRF (Mlawer et al., 1997; Iacono et al., 2008). It is designed to be accurate and computationally efficient, and has also been implemented in global circulation models. Ideally we would like to use as simple of a radiation scheme as possible, as they are very computationally intensive. Computational time can be saved, without necessarily sacrificing needed accuracy, in two ways: 1) by using a simpler (and therefore more efficient) scheme and 2) calling the radiation scheme to update conditions less frequently. We address the first point through the use of the RRTMG. To address the second point, we adjust the frequency with which the RRTMG is called as is discussed in Section C: Radiation time-step.

The most important physics packages for our simulations handle the treatment of the lowest part of the atmosphere, and its interaction with the underlying ocean surface. Rather than repeating the work of others, we refer to the published literature for guidance on our model physics. Both boundary layer parameterizations that we tested are based on the Mellor-Yamada-Janjic (MYJ) second-order boundary layer scheme and its associated surface layer scheme, which is based on Monin-Obukhov similarity theory (Janjic, 2002). An alternative formulation, the Mellor-Yamada-Nakanishi-Niino (MYNN; Nakanishi and Niino, 2004), was also included in our sensitivity tests. The details of how the schemes work are based on well-established theory, but their complexity exceeds the scope of this report.

The results of Otkin and Greenwald (2008) suggest that the MYJ scheme outperforms other schemes, including the Yonsei University (YSU) scheme (Hong and Noh, 2006). However, recent sensitivity tests performed by AER for a CO₂ emission study in an urban area showed little systematic difference between the two schemes (Nehrkorn et al., 2011). The role of the boundary layer parameterization was the main focus of our sensitivity tests in Task 1.

The MYNN and MYJ are based on the same equation that predicts turbulent kinetic energy (a variable that is required for this study). The ability of both parameterizations to predict turbulent kinetic energy is the reason that they were selected over the other options. The two schemes differ in the way that they calculate the mixing lengths (i.e. the distance over which turbulence mixes the atmosphere). The MYJ was developed to match observations whereas the MYNN was developed based on the results of idealized, high-resolution computer simulations known as large-eddy simulations (LES). The proponents of the MYNN claim that it is preferable, as it was developed to take into account a broad range of scales without the complexities induced by real-world conditions. However, it is unclear if this is a positive or negative for our applications to real-world conditions. Unlike the MYJ, the MYNN has no dependence on boundary layer height and does partially take into account the role of water vapor condensation. Details on the formulations of the MYJ and MYNN are available in Janjic (2002) and Nakanishi and Niino (2004) respectively.

The microphysics and cumulus schemes are known to greatly impact model forecasts of moist convection and remain an active area of research at AER (Adams-

Selin et al., 2011). For representation of processes associated with water vapor, we have identified the Ferrier cloud microphysics scheme as a compromise between efficiency and realism (Ferrier et al., 2002). It is more realistic than the simple warm rain scheme that neglects ice, but it is substantially faster than more sophisticated schemes that require multiple species of hydrometeors (e.g. droplets, snow, ice, etc). Because of its efficiency, the Ferrier scheme is used by NCEP in the operational WRF-NMM.

Given the wide range of scales in our simulations ranging from the tens of kilometers in the NARR to just hundreds of meters in our innermost domain, we will use a mix of cumulus schemes. For the outermost domain (horizontal resolution = 13.5 km), a parameterization will be needed to represent cumulus convection. The Kain – Fritsch scheme (Kain, 2004) is chosen as a compromise between the more computationally expensive Grell-Devenyi ensemble scheme (Grell and Devenyi, 2002), and the less realistic Betts-Miller-Janjic adjustment scheme (Janjic, 2000). The inner domains have high enough resolution to represent moist convective processes explicitly and in those domains no cumulus scheme will be used.

Running WRF

The WRF simulations for Task 1 were conducted on an SGI Altix ICE cluster with 2,048 processor cores using Intel compilers and SGI's MPT for message passing. In our experience, MPT has benchmarked about 8% more efficient than the comparable Intel and open-source message passing libraries. In order to limit the overhead associated with running WRF simulations in parallel, each simulation was run on 64 cores, allowing many simulations to be run concurrently. Excluding pre- and post-processing, a single model simulation of 30 hours took approximately 220 core-hours (about 3.5 hours of elapsed wall clock time running on 64 cores). Additional time, though not comparatively computationally intensive, is necessary to download and pre-process the data to initialize the WRF simulations and then post-process the data to extract the information of interest.

Choice of testing period

To ensure that the period simulated for our sensitivity testing was representative of a full climatology, we compared the wind statistics using the NARR climatology. In order to most faithfully represent the actual variability expected in a full climatology, we sought days that were independent samples (i.e., far enough apart to not be correlated) and not preselected to meet some preconception. Working under the assumption that more recent data will be of higher quality, we chose to sample from calendar year 2010. We don't expect that choosing a different year to test would impact our conclusions as the differences documented in this report show a diverse range of errors that show no seasonal patterns that would be indicative of the errors being caused by weather patterns that vary significantly from year to year. Our choice of starting date was decided by the distribution that most closely resembled that year's full climatology. The sampling strategy is the same as that which we plan to use to conduct our simulations in Task 2.

After inspecting the options, we settled on a subsample obtained by choosing every eleventh day starting January 7, 2010. Wind roses for the full year and this

subsample are shown in Figure 4 and Figure 5 respectively. The subsample adequately represents a maximum occurrence of northwesterly wind, often associated with cold frontal passage, and a secondary maximum of southerly winds, often associated with prefrontal warm air advection ahead of approaching frontal cyclones (like Nor'easters). The lack of easterly winds is typical of mid-latitudes. Note that a northwesterly wind is defined as a wind blowing from the northwest toward the southeast.

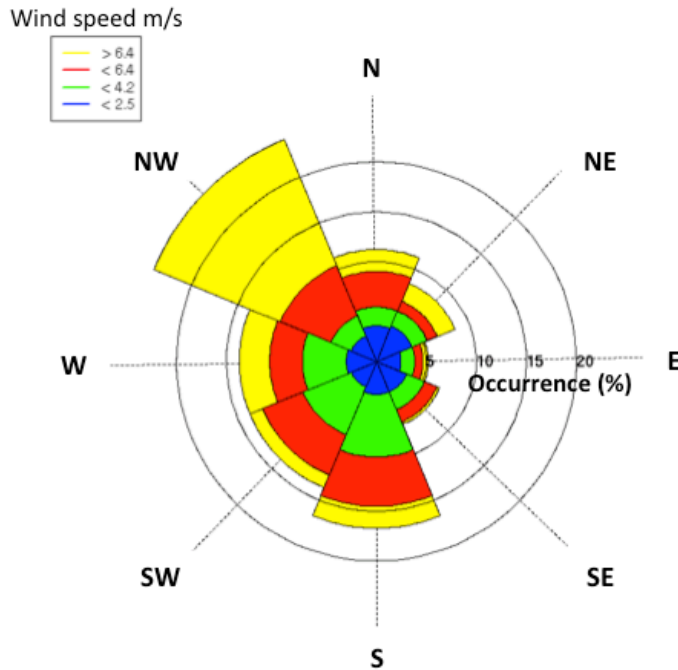


Figure 4: Wind rose for the OCS from the NARR historical reanalysis for the full calendar year 2010. Each slice is representative of a directional bin (labeled in degrees from north). The length of the slice is representative of the relative frequency (%), labeled along the x-axis. The colored portions correspond to a speed (m/s) bin for each direction (i.e., green is 2.5 - 4.2 m/s).




Figure 5: Wind rose for OCS from the NARR historical reanalysis for the dates used in our sensitivity testing: every eleventh day in 2010 starting on January 7. Each slice is representative of a directional bin (labeled in degrees from north). The length of the slice is representative of the relative frequency (%), labeled along the x-axis. The colored portions correspond to a speed (m/s) bin for each direction (i.e., green is 2.5 - 4.2 m/s).

We define ‘extreme’ events as those strong wind events that occur less frequently than once per year. In this fashion ‘extreme’ events are not likely to be captured by a sampling strategy that models a year or less of cases. For the same definitional reasons they will not contribute to the full climatology. If an event occurs often enough to be a significant contributor to the climatology it is by definition not extreme and will likely be captured by a ‘random’ sampling strategy.

For example, Nor’Easters are not defined as extreme as they occur several times per year. On the other hand, hurricanes are considered to be the dominant extreme events in the region as they occur less than once per year but are known to bring the strongest recorded winds to the region. Hurricanes’ impacts and frequency will be studied separately as part of the scope for Task 2.

On the other end of the spectrum, extremely quiescent conditions are ignored as they will not impact wind energy estimates; it makes no difference whether winds are light (well below cut-in wind speed for an offshore wind turbine) or completely calm. For example, whether we sample winds of 1 m/s or 0 m/s, neither will contribute to the wind energy estimates as both are below the cut-in speed at which turbines begin generating electricity. Insofar as these quiet days are

captured by random sampling they will be included, but no extra effort will be made to account for them. Such quiescent days are indicated by the blue pieces of the wind rose in Figure 4 and Figure 5.

Validation

Validation of numerical weather simulations over the oceans is complicated by the fact that there is typically very little observational data. For example, there are only two buoys near the OCS study area (see Figure 1) with only the southern buoy available in recent years. Additionally, buoys only provide wind measurements at an anemometer height of 5 m above the ocean surface. Even in calm seas this would be a substantial limitation in validating wind forecasts throughout the depth of the boundary layer. It is not at all clear how to interpret a 5 m wind observation coming on top of swells of comparable height. It is for this reason that we shifted our sensitivity test domain onshore (Figure 2).

The model domain for Task 1 was not shifted onshore to deemphasize the buoy data. To the contrary, the buoy was carefully left within the shifted model domain. Instead the domain was shifted to provide additional onshore observations for validation while still retaining the buoy data. Having the majority of the model domain offshore is desirable in Task 2 as it is the OCS that is the focus of the study. However, in Task 1 where validation and calibration is the primary goal, there is little purpose to include large offshore regions that are void of verification data. Instead offshore verification will rely heavily on the available buoy data.

Even with an onshore model domain, we still want more validation observations than those available through NWS's traditional Automated Surface Observing System (ASOS). To provide access to a greater number of observations we collected archived verification data from the Meteorological Assimilation Data Ingest System (MADIS) at NWS. The MADIS data is collected from a number of sources inside and outside traditional observing networks, including research networks and those run by state departments of transportation to monitor road conditions. For example, the Delaware Environmental Observing System (DEOS) run by the University of Delaware provides 30 stations. A map of observing stations used to verify our inner domain is given in Figure 6. Note that most observations are located in coastal regions, from both land and buoy observations.




Figure 6: Location of surface observations used to verify our sensitivity tests. MADIS observations are marked with circles. NWS stations are marked with triangles. Buoy 44009 is marked with the cross. The OCS area of interest is marked with the dashed lines and is the same as that in Figure 1.

Note that not all of the MADIS data is freely available as several data sets have usage and redistribution restrictions³. Only those observations that are freely available were used in this study. MADIS takes a lot of the overhead out of verification by streamlining collection, quality control, and distribution into one system. As part of the assimilation system, all surface observations are ostensibly delivered at observation heights of 2 m for temperature and humidity and 10 m for wind speed and direction.

Even with WRF simulations and MADIS data for verification in hand, it is not a trivial task to reconcile the forecasts with observations. For this task we employed

³ MADIS data and restrictions are available online at <http://madis.noaa.gov>

the Model Evaluation Tools (MET)⁴ verification software package from the NCAR Development Testbed Center. The MET software is maintained and developed with the support of the U.S. Air Force Weather Agency and National Oceanographic and Atmospheric Administration. We used the MET to generate “matched pair records” (i.e. pairs of corresponding forecasts and observations) for validation. The MET ingest post-processed forecasts and MADIS data to find pairs of forecasts and observations that are collocated in both time and space. To provide a manageable set we passed hourly forecasts from WRF to MET, and MET matched those hourly forecasts with all MADIS observations that were made within ± 1 hour from the forecast time. After checking for quality control, the matched pair records are output along with the relevant time and location information. Automated scripts calling MET were run on an IBM p1350 Beowulf cluster locally at AER, using approximately 24 hours of wall clock time.

⁴ More information on MET is available online at:
<http://www.dtcenter.org/met/users/index.php>.

C. Test results

Our sensitivity tests have revealed a set of model parameterizations that provide satisfactory model performance at an increased efficiency to allow for a greater sample of days to be modeled in Task 2. The update time interval for the radiation scheme was found to be less critical over the OCS than is generally required over land. The effect of the damping layer at the upper boundary of the model was found to be negligible for our purposes simulating wind speeds near the ocean surface. Both boundary layer parameterizations that we tested were found to perform equally well with no significant biases (i.e. the biases were of a lower magnitude than the measurement error of the observations used to verify them). The error statistics of the boundary layer parameterizations were found to be statistically indistinguishable and so the more efficient scheme is recommended.

Radiation time-step

As a general rule of thumb, the official WRF documentation recommends calling the radiation scheme in an interval in minutes that matches the horizontal grid resolution in km (i.e., for 10 km grid spacing the radiation scheme should be called every 10 minutes). For our innermost domain with a grid spacing of 500 m that translates to a recommend radiation update interval of 30 s. This update frequency is largely driven by a desire to keep the radiation balance in line with clouds passing over complex terrain. In such an environment, the ground can heat up and cool down relatively rapidly and induce small-scale circulation changes. For our applications, the ocean surface is relatively flat and responds very, very slowly to changes in the radiation balance as, unlike land, the ocean surface is mixed by waves and the wind. Therefore we expect that calling the radiation scheme less frequently to update conditions will have little impact on the results.

To test the need for such rapidly changing radiation conditions we ran a sample test case changing only the rate at which the radiation was updated. We compared a high-rate interval, set to 1 min for our test, to what we consider a more reasonable interval, 10 min. An example wind field from the high-rate radiation run is shown in Figure 7. Our results showed that even after 30 hours of integration, the model differences of wind speed were small and represented slight phasing differences in the passage of wind gusts (Figure 8). The total wind resources were very similar between the two tests. The total differences in wind speed were on the order of 6% and represented slight differences in the passage of transient gusts. There was no significant bias in either direction so little impact is expected in the final results.




Figure 7: Surface (10 m) wind speed from a 30-hour simulation initialized at 00 UTC on March 1, 2010 with 1 minute radiation updates.

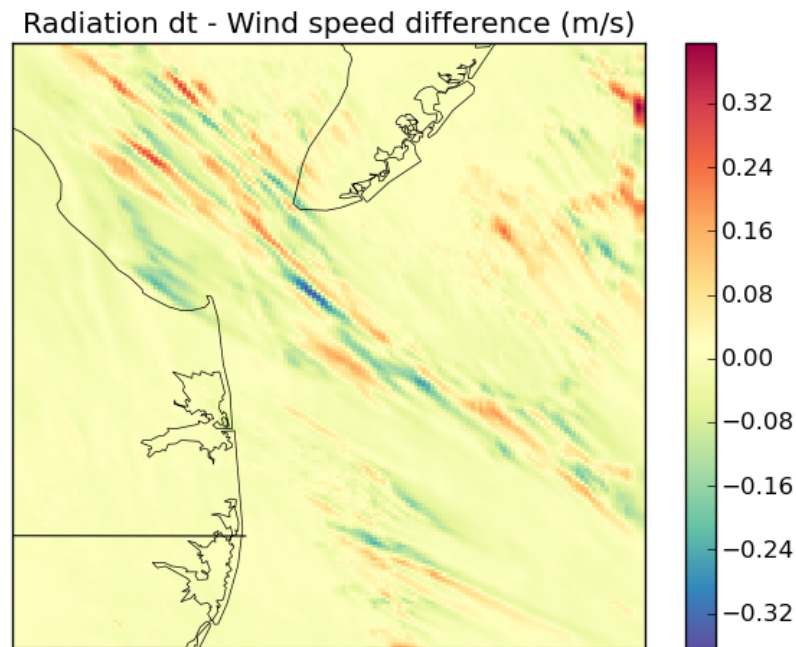


Figure 8: Difference between surface (10 m) wind speed from 1 minute and 10 minute radiation updates from a 30-hour simulation initialized at 00 UTC on March 1, 2010. Note that the differences of up to 0.32 m/s correspond to roughly a 6% difference and are largely related to the positioning of wind gusts associated with turbulent mixing in the boundary layer.

Damping layer

In order to prevent the contamination of forecasts from kinetic energy (in the form of “gravity waves”) reflected off of the upper boundary, operational weather forecast models are run with a layer at the top of the model that absorbs kinetic energy. This differs from traditional analytical solutions and laboratory experiments, which typically have a rigid lid that reflects incoming kinetic energy. The handling of the upper boundary condition makes a large difference in the flows that develop near the top of the model, especially in the stratosphere. Manifestations at the surface are less common during short time scales, though they are often seen in the development of down-slope windstorms in the lee of mountains. Damping layers are most important in simulations with large sources of upward propagating kinetic energy like deep convection and flow over mountainous terrain. To test for a need for a damping layer in our retrospective forecasts, we repeated the same sensitivity test used in the radiation test (sticking with the less frequent 10 minute radiation updates), but this time testing either no damping layer or the state-of-the-science damping layer of Klemp et al. (2008). The results of the damping layer test are similar to that of the radiation test. The differences were all

under 0.5 m/s and representative of slight differences in the time of gusts (Figure 9).

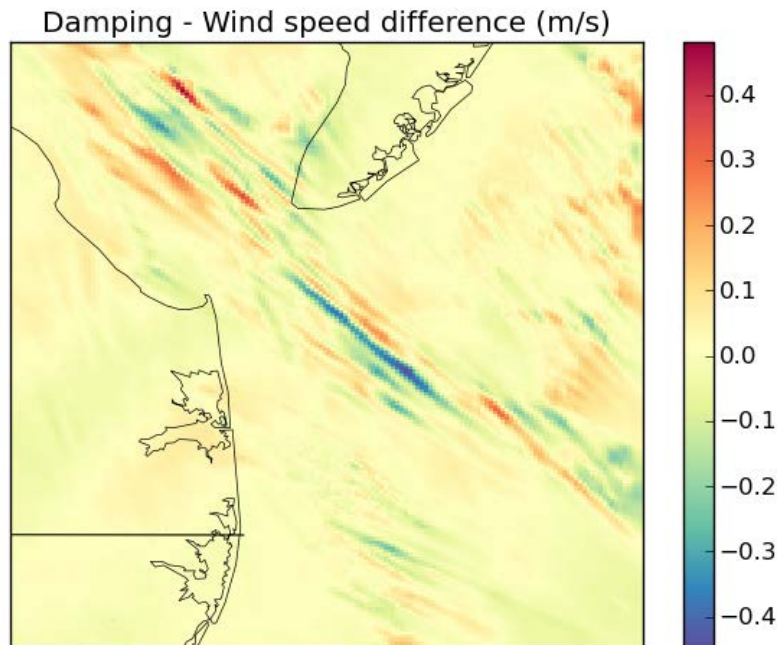


Figure 9: Difference between surface (10 m) wind speed from simulation with and without an upper damping layer from a 30-hour simulation initialized at 00 UTC on March 1, 2010. Note that the differences are largely related to the positioning of small-scale features and show no systematic bias.

Boundary layer parameterization

As previously mentioned, the sensitivity tests were verified against observations collected by MADIS from a number of coastal locations (Figure 6), both ground stations and buoys. For an unknown reason, the buoy closest to the area of interest (44009) was not included in the MADIS data. It is possible that MADIS quality control is excluding the buoy as it started experiencing a loss of humidity data on 21 December 2010 and has been suffering additional problems with its ocean surface wave data. Accordingly that buoy was verified separately, which helped to underline some unexpected results that are explained later. The dominantly coastal MADIS observations showed the model forecasts to have a clear and consistent high bias (i.e., predicted wind speeds were on average higher than observed) while the observations from buoy 44009 showed no consistent bias.

Figure 10 and Figure 11 show scatter diagrams of model error versus the forecast wind speed for the MYJ and MYNN boundary layer parameterizations respectively. All simulations were conducted with 10 minute radiation updates and no damping. Both parameterizations show a clear positive bias (i.e. systematic over-prediction) that increases with increasing wind speed. Each point corresponds to a matched pair record for an hourly forecast and MADIS observation. The sharp upper

cutoff is a physical limitation as points above that line would imply that the error (forecast - observation) is greater than the observed wind speed. This requires a negative wind speed, which is impossible. This bias shows a clear limitation to the forecasts at the verification sites, but does little to explain the bias.




Figure 10: Model error (forecast - observation) versus the forecast wind speed (m/s) for all forecasts using the MYJ boundary layer parameterization. The blue line is the least-squares trend indicating that the wind speed forecasts appear to be roughly double the observations.

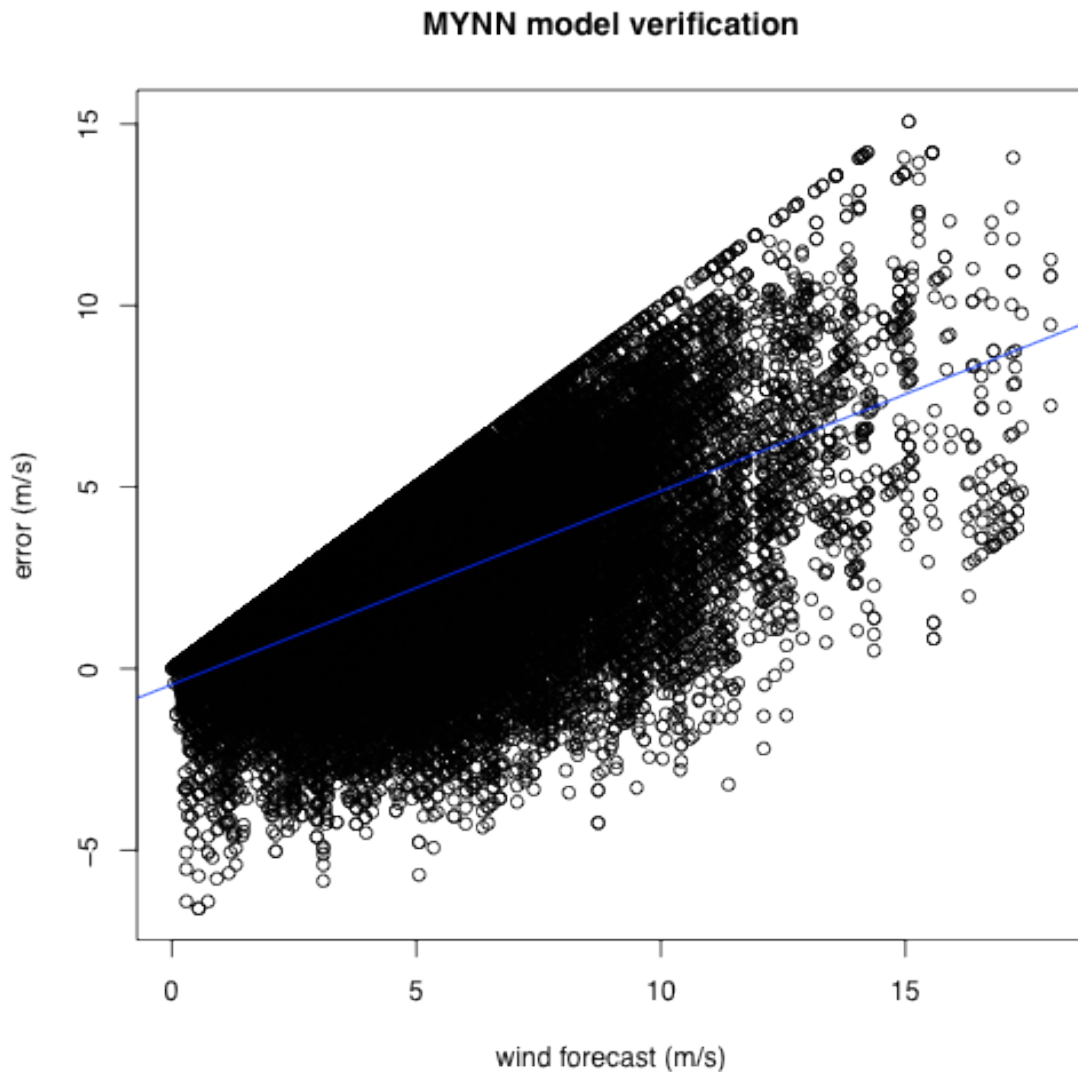


Figure 11: Model error (forecast - observation) versus the forecast wind speed (m/s) for all forecasts using the MYNN boundary layer parameterization. The blue line is the least-squares trend indicating that the wind speed forecasts appear to be roughly double the observations.

The error statistics for the points shown in Figure 10 and Figure 11 are summarized in Table 2. Disconcertingly the verification statistics show biases that are so large that it is hard to see how they could be useful to create a resource estimate without a statistical bias correction. The bias seems to show no seasonal correlation (Figure 12, Figure 13) meaning that it is unlikely that the bias is due to any one weather pattern. This bias is consistent with another set of simulations for August 2010 on a 1.33 km grid around Boston using the MYJ scheme that are currently being analyzed at AER. Those Boston simulations show systematic biases from +2 m/s up to +6 m/s compared to MADIS observations. For the Boston simulations, comparison against standard NWS observations showed no clear bias like that seen in the MADIS verification, suggesting a systematic underestimation of wind speed in the non-

standard MADIS observations. More details on the Boston simulations are given in Appendix: Boston simulations.

Table 2: Root mean square error $RMSE = (\text{mean}((\text{forecast} - \text{observed})^2))^{1/2}$, and mean error -- bias = $\text{mean}(\text{forecast} - \text{observed})$, for both boundary layer parameterizations verified against MADIS observations, buoy 44009, and official surface stations at Georgetown, DE (KGED), Wildwood, NJ (KWWD), and Ocean City, MD (KAXB). The values given for buoy 44009 included the actual observations at 5 m height and those "corrected" to 10 m height in parentheses.

| <u>Case:</u> | <u>MYI</u> | | <u>MYNN</u> | |
|---------------------|-------------------|-------------------|--------------------|-------------------|
| | RMSE (m/s) | Bias (m/s) | RMSE (m/s) | Bias (m/s) |
| MADIS | 3.53 | 2.68 | 3.36 | 2.37 |
| KGED | 3.13 | 0.86 | 2.90 | 0.37 |
| KWWD | 4.35 | 1.46 | 4.01 | 1.02 |
| KAXB | 3.54 | 1.47 | 3.14 | 0.81 |
| Buoy 44009 | 2.24 (2.37) | 0.17 (-0.38) | 2.33 (2.41) | 0.25 (-0.30) |

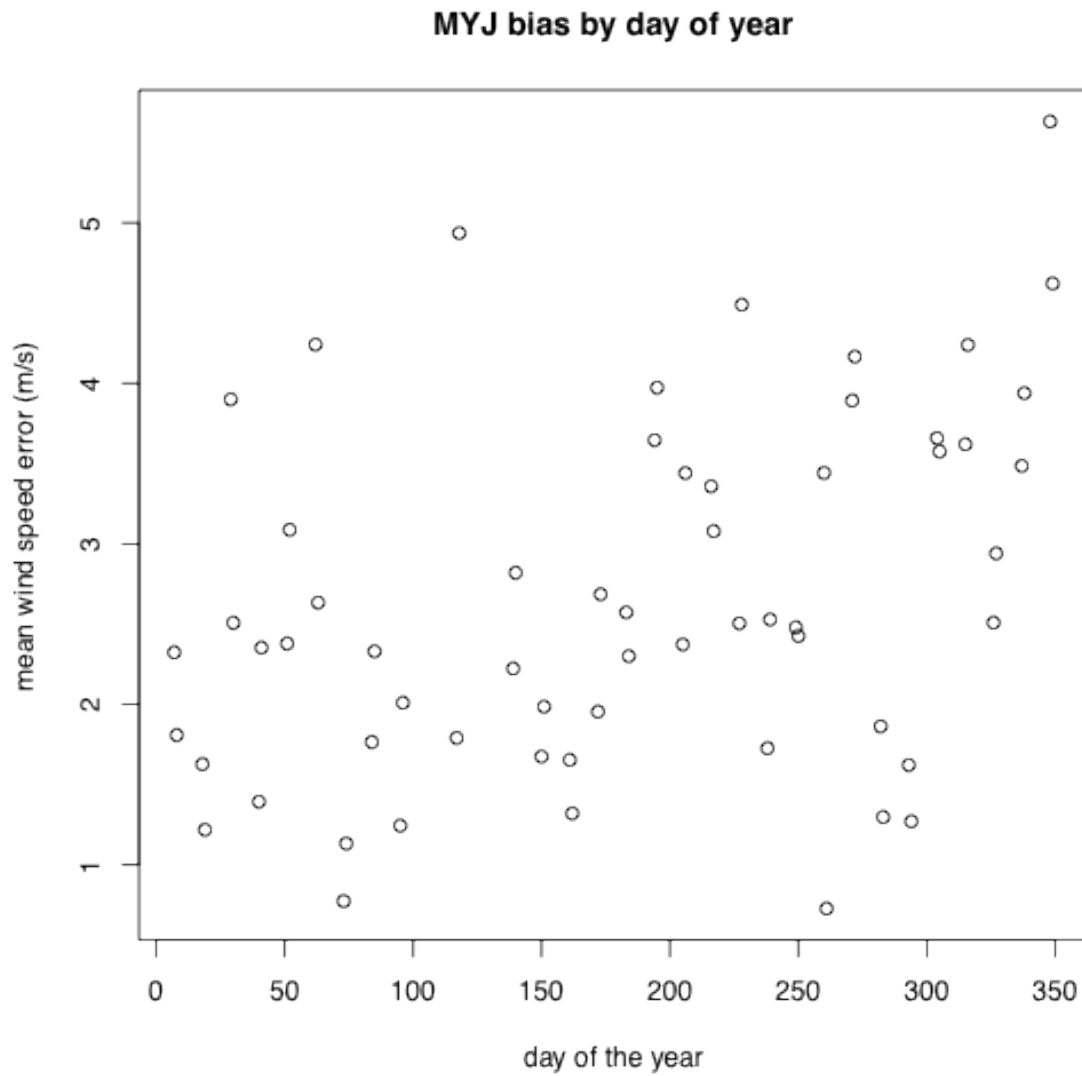


Figure 12: The mean wind speed error (forecast - MADIS observations) for each day 11th day of the calendar year 2010 (i.e. those modeled) with the MYJ boundary layer parameterization.

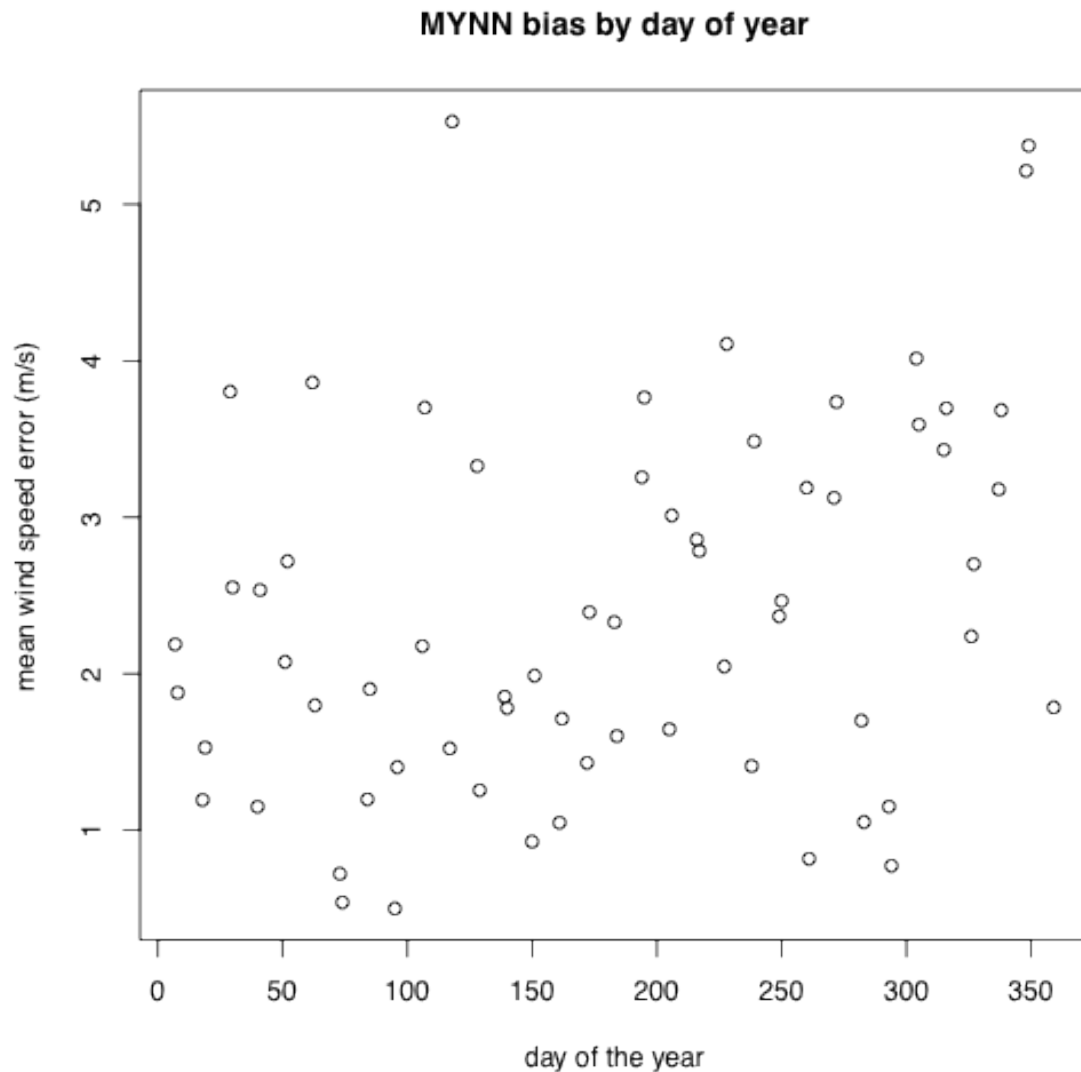


Figure 13: The mean wind speed error (forecast - MADIS observations) for each day 11th day of the calendar year 2010 (i.e. those modeled) with the MYNN boundary layer parameterization.

Unlike in our Boston simulations, there seems to be no diurnal trend in the bias for our MYJ or MYNN simulations (Figure 14, Figure 15), whereas the Boston simulations showed a bias minimum before dawn and maximum in the afternoon. The average forecast and observed values show a very similar shape by hour of the day with a consistent bias showing little systematic variation. This is an important point as our 30-hour integrations cover the time period of 00 - 06 UTC twice in each simulation, once at the beginning and once at the end of each 30 hour period. This would provide double the weighting to the hours 00-06 UTC compared to the rest of the day when calculating simple error statistics, so it could have been an obvious source of a systematic bias.

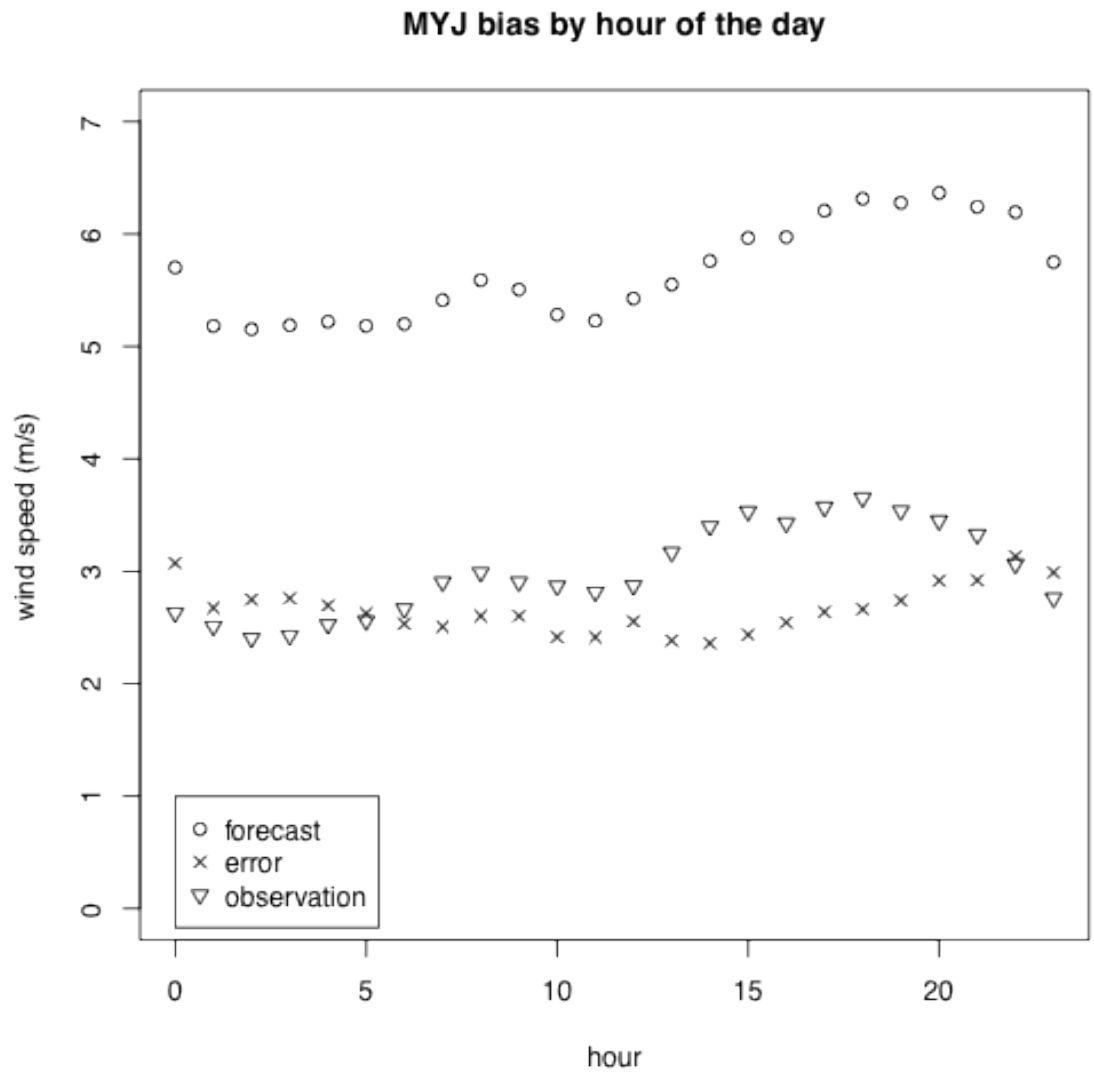


Figure 14: The average forecast, error, and observed wind speed by hour of the day averaged across all simulations with the MYJ boundary layer parameterization.

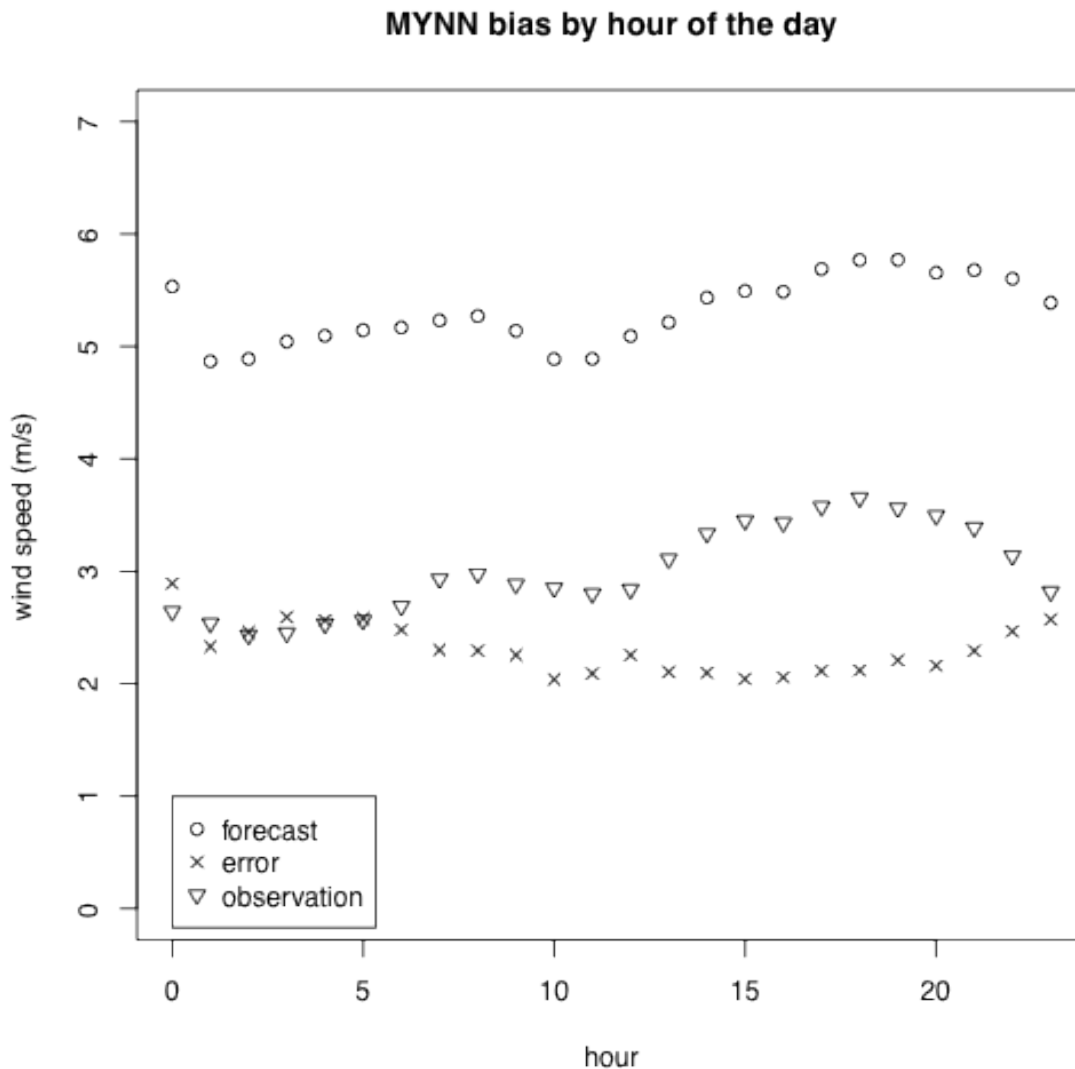


Figure 15: The average forecast, error, and observed wind speed by hour of the day averaged across all simulations with the MYNN boundary layer parameterization.

The systematic high bias is also obvious in the histograms of the forecasts (Figure 16, Figure 17) compared to those of the MADIS observations (Figure 18). The shape of the observation curve shows a clear maximum at the lowest wind speeds, while the forecast calls for a distribution with a maximum around 5 m/s. The model errors (forecast – observed) show a distribution with a single mode, but clearly shifted toward a high bias of about 2 m/s (Figure 19, Figure 20). If there were no systematic bias, we would expect a normal distribution centered on zero.




Figure 16: Histogram of model forecasts using the MYJ boundary layer parameterization from the MADIS observation locations. The histogram was built using a bootstrapping method sample equally from all hours of the days to remove a sample bias inherent in 30-hour forecasts.

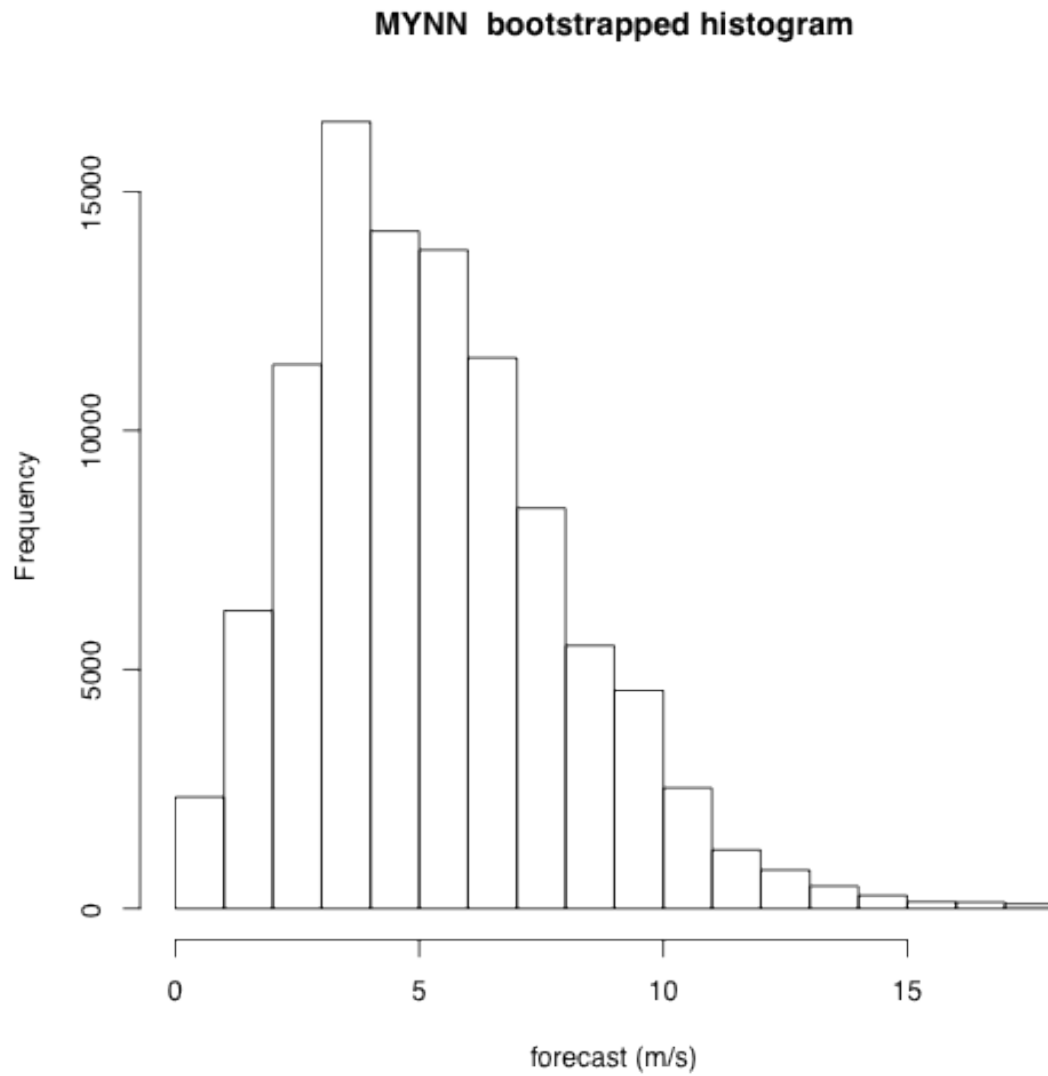


Figure 17: Histogram of model forecasts using the MYNN boundary layer parameterization from the MADIS observation locations. The histogram was built using a bootstrapping method sample equally from all hours of the days to remove a sample bias inherent in 30-hour forecasts.




Figure 18: Histogram of MADIS observations taken from the MYJ match pair records. The distribution from the MYNN match pair records is nearly identical as the observations are drawn from the same data so only one distribution is presented. The histogram was built using a bootstrapping method sample equally from all hours of the days to remove a sample bias inherent in 30-hour forecasts. Note the clear maximum at the lowest wind speeds.

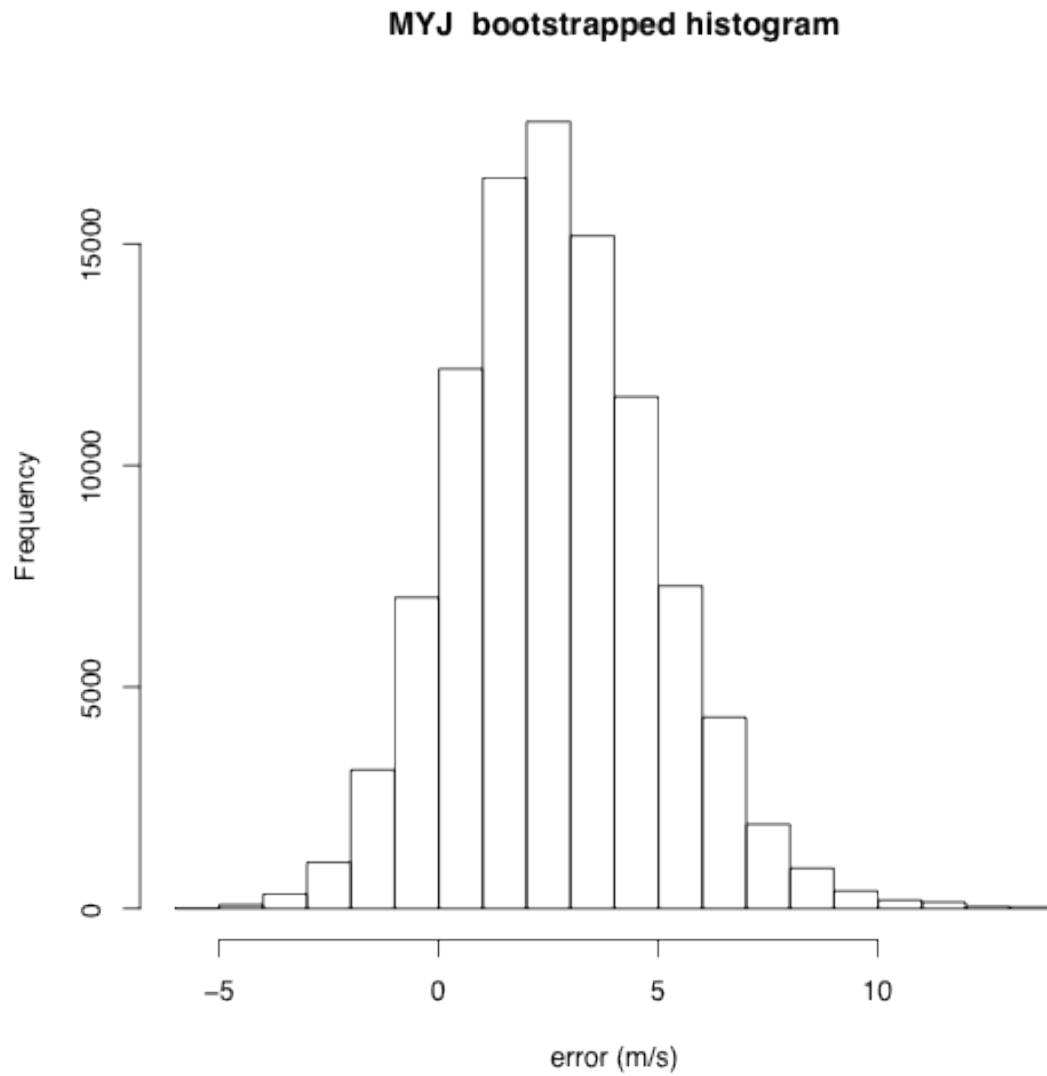


Figure 19: Histogram of model forecast errors from the MYJ simulations. The histogram was built using a bootstrapping method sample equally from all hours of the days to remove a sample bias inherent in 30-hour forecasts. Note the maximum is shifted to a 2-3 m/s bias.




Figure 20: Histogram of model forecast errors from the MYNN simulations. The histogram was built using a bootstrapping method sample equally from all hours of the days to remove a sample bias inherent in 30-hour forecasts. Note the maximum is shifted to a 2-3 m/s bias.

While the coastal MADIS observations suggest that the model has a significant, but consistent, high bias, the same cannot be said of the offshore observations. Surface wind speed observations from buoy 44009 show a wind speed distribution that much more closely resembles the shape of the distribution of model predictions (Figure 21). The error statistics (forecast-observations) show a bell-shaped distribution (Figure 22, Figure 23) with a bias that is comparable to the wind speed measurement resolution (± 0.1 m/s) and well below the measurement accuracy (± 1 m/s) for the ARES instrument package aboard that specific buoy (Table 2). Comparison with buoy data is complicated by the fact that the buoy surface wind observations are taken at an anemometer height of 5 m rather than the standard surface observation height of 10 m. The statistics for 10 m winds shown in parentheses in Table 1 are computed by correcting the 5 m wind speed using a

method recommended by the National Data Buoy Center⁵, the power law correction of Hsu et al. (1994), $u_2 = u_1(z_2/z_1)^P$, where u_2 is the wind speed adjusted to $z_2 = 10$ m, u_1 is wind speed measured at $z_1 = 5$ m, and $P = 0.11$ is an empirically-derived coefficient for the ocean. This correction assumes neutral atmospheric stability, a flat sea surface, and is derived empirically from oil platform and buoy data. Given that the sea surface is often not flat and that a small buoy will be pitching about in the waves, it is not obvious how to correct for any biases in the buoy data so it is likely wise to stick with the actual surface observations, which show a smaller bias and root mean square error than the 'corrected' values. In this case we are comparing model forecasts at 10 m to buoy observations at 5 m.

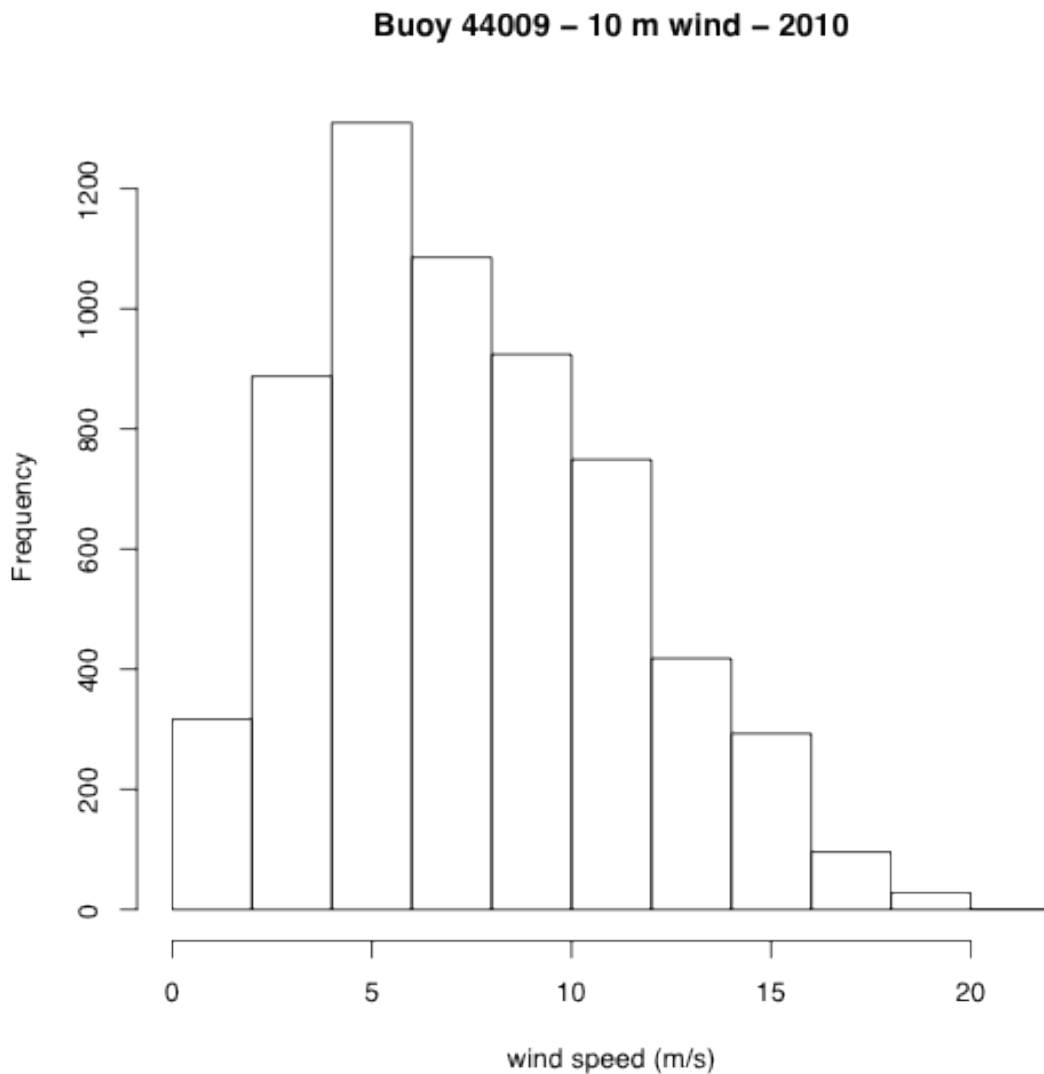


Figure 21: Histogram of all surface wind speed observations at buoy 44009 in calendar year 2010.

⁵ Available online at http://www.ndbc.noaa.gov/adjust_wind.shtml

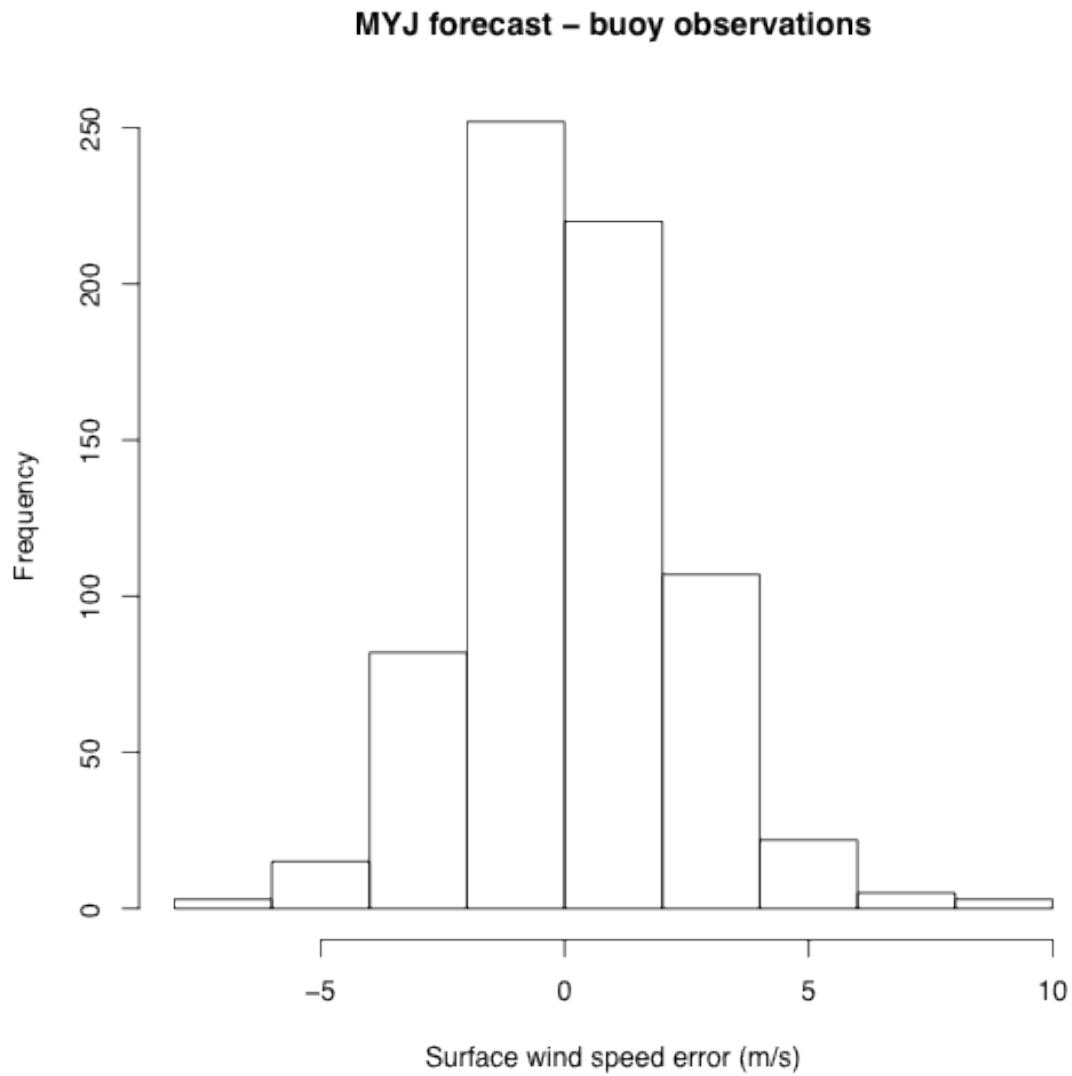


Figure 22: Histogram of wind speed error (forecast – observation) at buoy 44009 for simulation using the MYJ boundary layer parameterization.

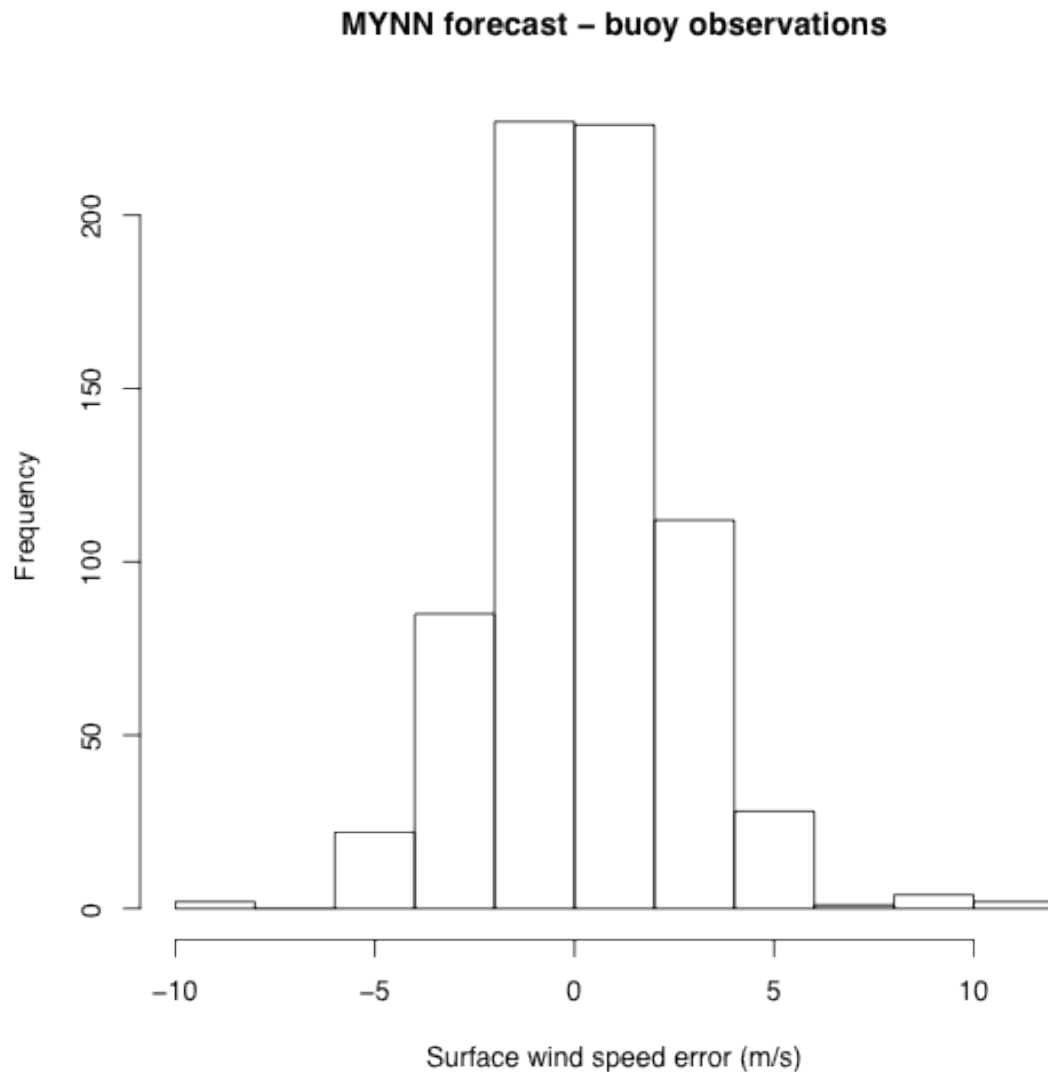


Figure 23: Histogram of wind speed error (forecast – observation) at buoy 44009 for simulation using the MYNN boundary layer parameterization.

To test whether the differences between MYJ and MYNN in the error statistics shown in Table 2 are statistically significant, we apply Fisher’s F-test and Welch’s two-sample t-test. The F-test has one tail so we are looking for p-values over 0.05 to allow us to assume that the two distributions have the same variance. The t-test has two tails so we are looking for p-values over 0.025 and less than 0.975 to say that the means of the distributions are not statistically different. By comparing the error statistics of the model errors from the MYJ and MYNN boundary layer simulations, we see an F-test p-value = 0.508 and a t-test p-value = 0.508 indicating that the differences between the two error distributions are not statistically significant in either variance or mean.

The cause of the disagreement between the model forecasts and MADIS observations is not immediately clear. The bias-free agreement between buoy 44009 and the model predictions argues against us needing to concern ourselves

with this discrepancy for the purpose of a climatology on the OCS. A detailed investigation is beyond the scope of this project, but a couple of possibilities come to mind. It is possible that the boundary layer parameterizations are simply too aggressive in mixing momentum down to the surface over land where stable boundary layers are more likely than over the ocean. The bimodal distributions of wind speed observations at two of the official NWS surface stations in the region of interest suggests that this is at least partially contributing to the disagreement (Figure 24, Figure 25). This is supported by the seemingly stronger secondary maximum at the lowest wind speeds at Georgetown, DE (KGED), which is further inland than the station at Wildwood, NJ (KWWD).

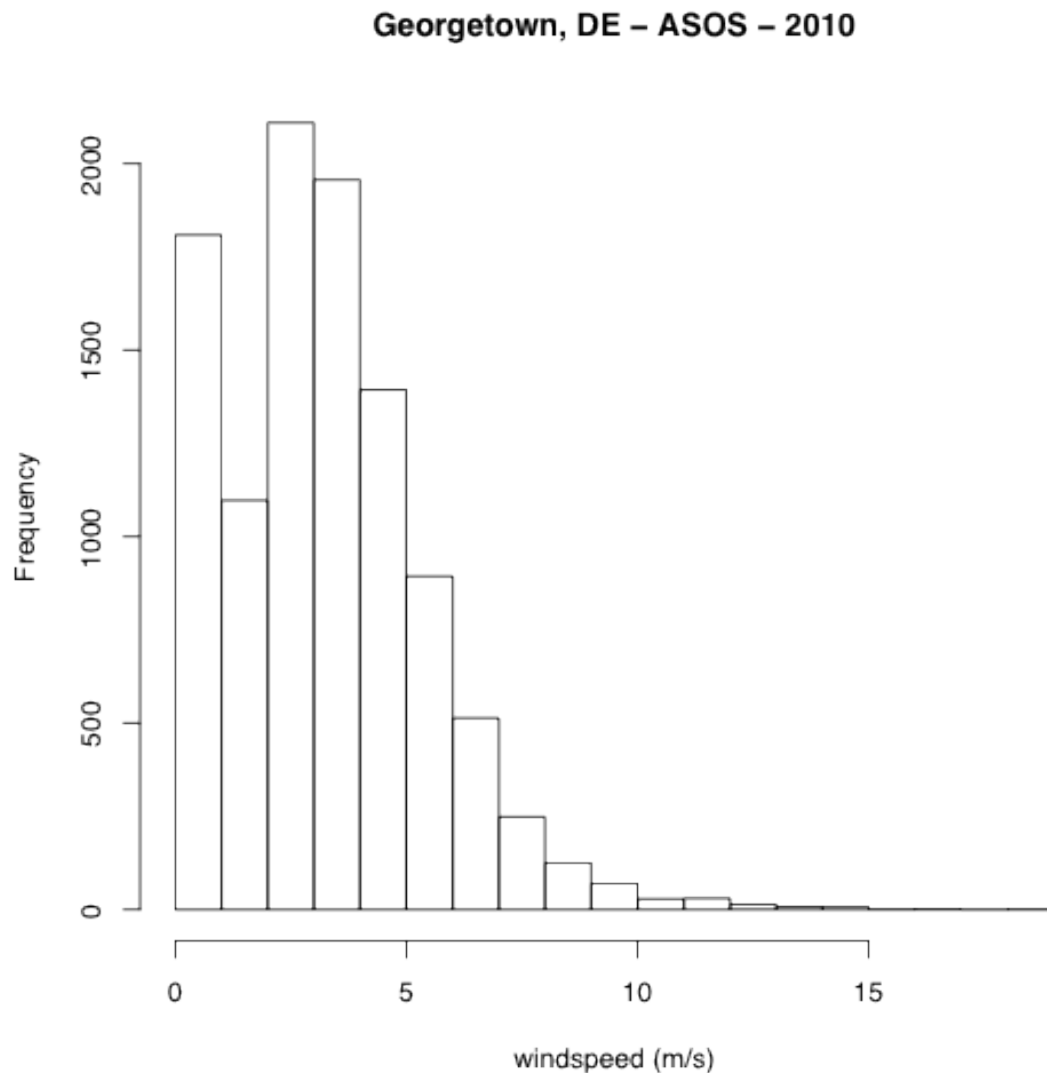


Figure 24: Histogram of all surface wind speed observations at Georgetown, DE (KGED) in calendar year 2010.

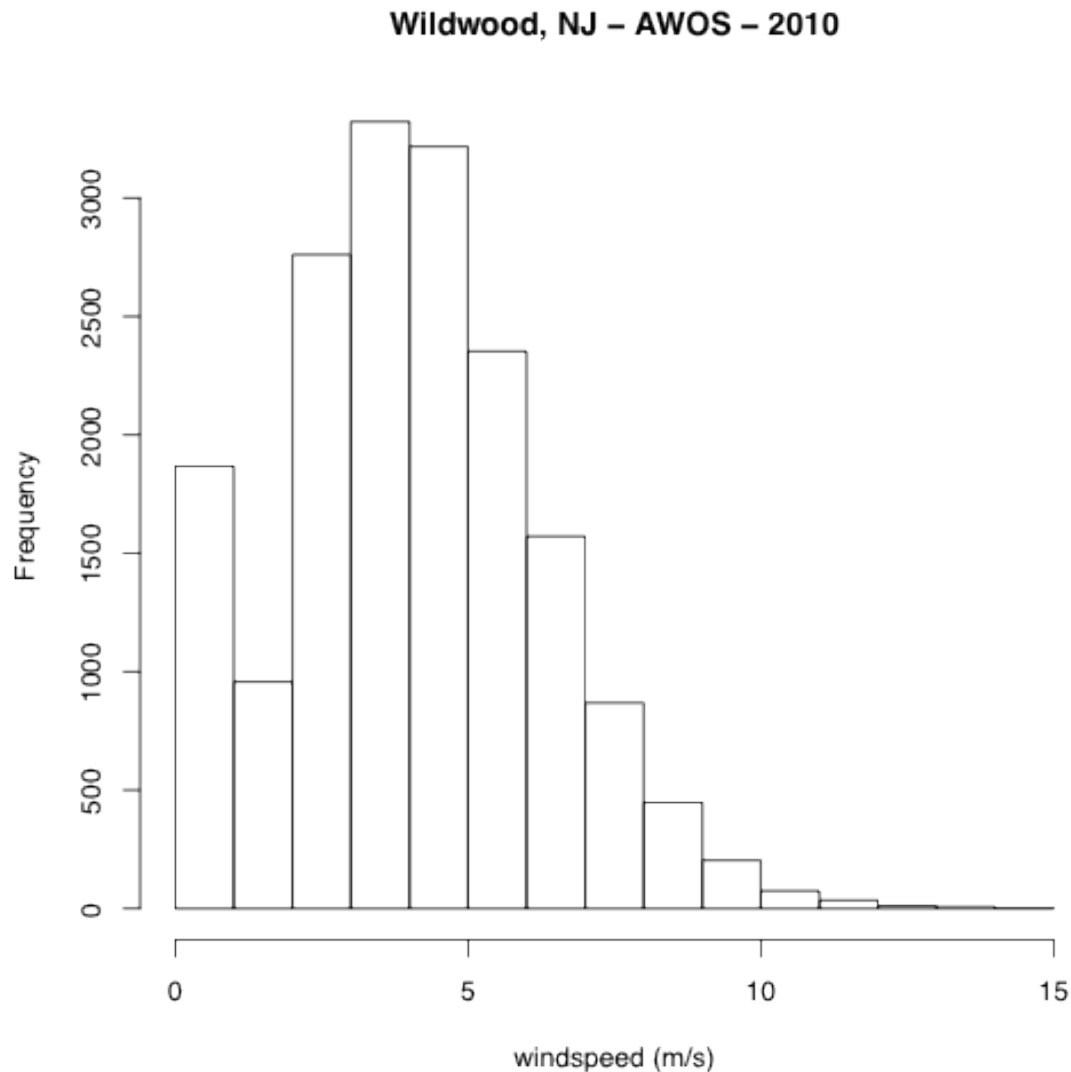


Figure 25: Histogram of all surface wind speed observations at Wildwood, NJ (KWWD) in calendar year 2010.

However, results from our Boston simulations suggest that a more important factor is the low-wind speed bias of the non-standard MADIS observations, which is not present in standard NWS surface wind observations. Most of the observations come from unofficial surface stations and are assimilated within MADIS. The observations are ostensibly all given at 10 m above ground level, but it is not clear that this is the height at which all measurements are taken. For example, the ASOS station at KGED, like all ASOS stations, is located in an exposed location with the anemometer at 10 m height (Figure 26). Further investigation would be required to determine how MADIS assimilates observations from stations like KGED with those from unofficial networks like DEOS, whose stations clearly do not measure 10 m winds (Figure 27). It appears many of the stations may instead be measuring 5 m (or lower) winds, much like buoy 44009 (Figure 28). The offshore buoy 44009 near the area of interest has verified well against our model simulations. The second

buoy in the area was not available for recent years and so no verification data was collected.




Figure 26: Photograph of the Georgetown, DE (KGED) ASOS station. Courtesy of the National Climatic Data Center.




Figure 27: The DEOS station in Ellendale, DE. Photo from the 17 January 2008 DEOS newsletter available at <http://www.deos.udel.edu/news/011708.html>.



Figure 28: A photo of a 3-meter discus buoy from NOAA courtesy of the National Data Buoy Center. The buoy shown is (probably) not buoy 44009, but rather an identically equipped unidentified buoy. The original image is available at http://www.ndbc.noaa.gov/images/stations/3m_mini.jpg.

WRF simulations using the MYJ and MYNN boundary layer parameterizations were verified against unofficial MADIS observations, three official ASOS station, and

an official buoy on the OCS. All runs were conducted using the NARR reanalysis and did not directly include surface observations. This allows us to use the surface observations as independent verification data. Results showed the MADIS observations to have significant data quality issues involving unrealistically low wind speed measurements. Results from three official onshore ASOS stations showed a smaller bias than the MADIS observations, but still indicated the model over predicts onshore wind speeds. However, buoy observations from on the OCS indicate the model has no statistically significant bias. The verification statistics for the MADIS, ASOS, and buoy observations are summarized in Table 2.

D. Recommendation

In summary, we recommend using the WRF model configurations as described in Table 3 and Table 1 with the MYJ boundary layer scheme, a radiation update frequency of 10 minutes, and an upper-level damping layer. The MYJ parameterization is slightly more computationally efficient which, when combined with the 10 minutes radiation updates and a damping layer to ensure stability, will allow for greater sample diversity for a given computational allowance.

This model configuration, in conjunction with the resampling strategy detailed in our proposal, will be capable of providing the output parameters required for Task 2. These requirements were specified in the contract document, and refined during the meeting at the start of the contract, as follows:

- Areal coverage: encompass the RFI
- Temporal coverage: for 30 years at 10-minute intervals
- Horizontal resolution: 500m horizontal grid resolution
- Vertical levels: 30, 60, 90, 120, 150, 210m MSL
- Atmospheric fields at vertical levels:
 - wind speed
 - wind direction
 - specific humidity
 - turbulent kinetic energy
 - air pressure and temperature
- Atmospheric fields at near-surface level:
 - 2m temperature and specific humidity
 - 10m u- and v-wind
 - surface sensible heat and moisture flux
 - friction velocity
 - sea-surface temperature
 - planetary boundary layer height

Validation of our sensitivity tests has revealed that we have identified a model configuration that reproduces the wind speed climatology with little bias on the OCS (compared to 1 offshore buoy measurement). However, the forecasts over land show a significant positive bias compared with observations collected using MADIS. Based on independent WRF model run verifications for our simulations and those from the Boston area, the main cause for this disagreement is a low wind-

speed bias present in the non-standard MADIS observations. A bias of this magnitude was not found when compared against standard NWS observations.

Over land, a second source of wind speed over-prediction is the inability of our model to forecast low wind speeds associated with a stable surface layer caused by radiational cooling of the ground at night. Over the ocean, where conditions are much more often neutrally stable (i.e. well mixed), this limitation is not a hindrance. As the entire area for data delivery at the end of Task 2 is on the OCS, we will ignore this limitation. However, this issue is mentioned as a cautionary tale to those seeking to apply our methodology onshore. Further study is necessary to quality control the validation data and test alternative model configurations for land-based wind resource modeling.

Our validation statistics provide us with little reason to choose either the MYJ or MYNN boundary layer parameterization over the other. Both sets of validation statistics showed biases that were similar in magnitude to the measurement precision and smaller than the measurement error. The F-test and t-test statistics showed the MYJ and MYNN error statistics are not statistically different from each other. Given that the error statistics from the two parameterizations are nearly identical, the two parameterizations are likely interchangeable for our purposes.

Based on the radiation sensitivity test we have determined that we can easily use less frequent radiation updates than is general recommended for running WRF – every 10 minutes instead of every 30 seconds in our case. The impact of this choice will result in slight phasing differences in the forecast, but will not impact the wind resource estimates. Likewise, the results appear insensitive to the application of an upper-level damping layer. While we did not apply one in the control sensitivity tests, we see little to lose in doing so in the operational runs. If anything, it will provide us with greater stability during periods of strong vertical motion (e.g. thunderstorms and lift associated with frontal passage) without significantly impacting the results.

Appendix

Location of vertical levels

In total we used 44 vertical levels, which correspond to 43 layers (WRF uses a staggered grid in the vertical, using layer midpoints for mass and horizontal velocity, and layer interface levels for vertical velocity). As WRF does not support vertical nesting (this becomes important when considering how the nests pass information back and forth), the vertical levels were the same between all of the telescoping nests. Vertical levels in WRF are manually specified using terrain-following “Eta coordinates”, η , defined as:

$$\eta = \frac{p - p_T}{p_S - p_T}$$

where p is the pressure at level η , and the subscript S and T refer to the model surface and top respectively. All simulations were computed using the 44 vertical levels listed in Table 3 and show in Figure 29.

Table 3: Vertical η levels used in the WRF simulations.

| | | | |
|--------|-------|-------|-------|
| 1 | 0.853 | 0.432 | 0.143 |
| 0.9939 | 0.826 | 0.398 | 0.125 |
| 0.9898 | 0.799 | 0.366 | 0.107 |
| 0.9858 | 0.772 | 0.336 | 0.091 |
| 0.9817 | 0.722 | 0.307 | 0.075 |
| 0.9757 | 0.674 | 0.279 | 0.061 |
| 0.9676 | 0.628 | 0.254 | 0.047 |
| 0.951 | 0.585 | 0.229 | 0.034 |
| 0.931 | 0.544 | 0.206 | 0.022 |
| 0.906 | 0.505 | 0.184 | 0.011 |
| 0.879 | 0.467 | 0.163 | 0 |




Figure 29: Distribution of vertical levels.

Model time step

To ensure that the simulation is stable, we must prevent violation of the Courant-Friedrichs-Lewy (CFL) condition (Courant et al., 1928). The CFL condition guarantees that the atmospheric flow will not travel more than one grid space in a single time-step. A general rule of thumb for WRF is to set the time-step (in seconds) equal to six times the horizontal resolution (in kilometers). For the inner nest with a horizontal resolution of $\frac{1}{2}$ km, that corresponds to a time step of 3 s. However, because we are dealing with such a fine resolution in the boundary layer, a shorter time step on the order of 2 s is needed to maintain numerical stability during periods of high turbulence or strong updrafts. Simulations on similar spatial grids have shown a need for time steps of 1 – 2 s for turbulent flows in areas of complex terrain (Woods and Smith, 2011).

Choice of sampling interval

The sampling period of every 11 days was selected based on the total computational allowance for our sensitivity tests. While it is a relatively coarse sampling interval compared to the time-scale of mid-latitude weather systems, it does a remarkably good job of reproducing the general characteristics of the annual wind climatology (Figure 4, Figure 5). The 11 day sampling interval results in 41 simulated days when accounting for the fact that each simulated day last 30 hours rather than 24 hours. This results in approximately 11% of the year being

simulated, which we consider a satisfactory sample. The wind roses for a 7-day and 4-day sampling interval are given in Figure 30 and Figure 31 respectively. Note that the general shape of the distributions is unchanged. The climatology is dominated by southerly and northwesterly winds associated with mid-latitude frontal cyclones.

As the 4-day interval has the most sampled days, it most closely resembles the full annual climatology. An exact match is not needed for Task 1 as the goal is simply a sensitivity test to calibrate our model configuration. As Task 2 will sample a longer time period (3 years) with the 11-day interval, the total number of sampled days will be very similar to the 4-day interval for a year. For example, every 4th day for a year will sample 91 days while every 11th day for three years will sample 99 days. Therefore we expect the wind rose for Task 2 to closely resemble Figure 31. That sample captures wind from every direction and closely resembles the full climatology.




Figure 30: Wind rose for OCS from the NARR historical reanalysis for the dates used in our sensitivity testing: every seventh day in 2010 starting on January 7. Each slice is representative of a directional bin (labeled in degrees from north). The length of the slice is representative of the relative frequency (%),

labeled along the x-axis. The colored portions correspond to a speed (m/s) bin for each direction labeled in the upper-left (i.e., green is 2.5 – 4.2 m/s).




Figure 31: Wind rose for OCS from the NARR historical reanalysis for the dates used in our sensitivity testing: every fourth day in 2010 starting on January 7. Each slice is representative of a directional bin (labeled in degrees from north). The length of the slice is representative of the relative frequency (%), labeled along the x-axis. The colored portions correspond to a speed (m/s) bin for each direction labeled in the upper-left (i.e., green is 2.5 – 4.2 m/s).

Boston simulations

A set of WRF simulations for a project that is independent of this study was conducted at AER concurrent with this study. The simulations used a very similar model configuration, including the MYJ boundary layer parameterization. Those simulations are herein referred to as the Boston simulations as the inner study domain with 1.33 km resolution is centered on southern New England. The Boston simulations were verified against a network of MADIS observations that were much more dense than the observations available to this study (Figure 32). The numerous

MADIS observations were overwhelmingly located over land. The few buoy observations that were included were limited to coastal regions.

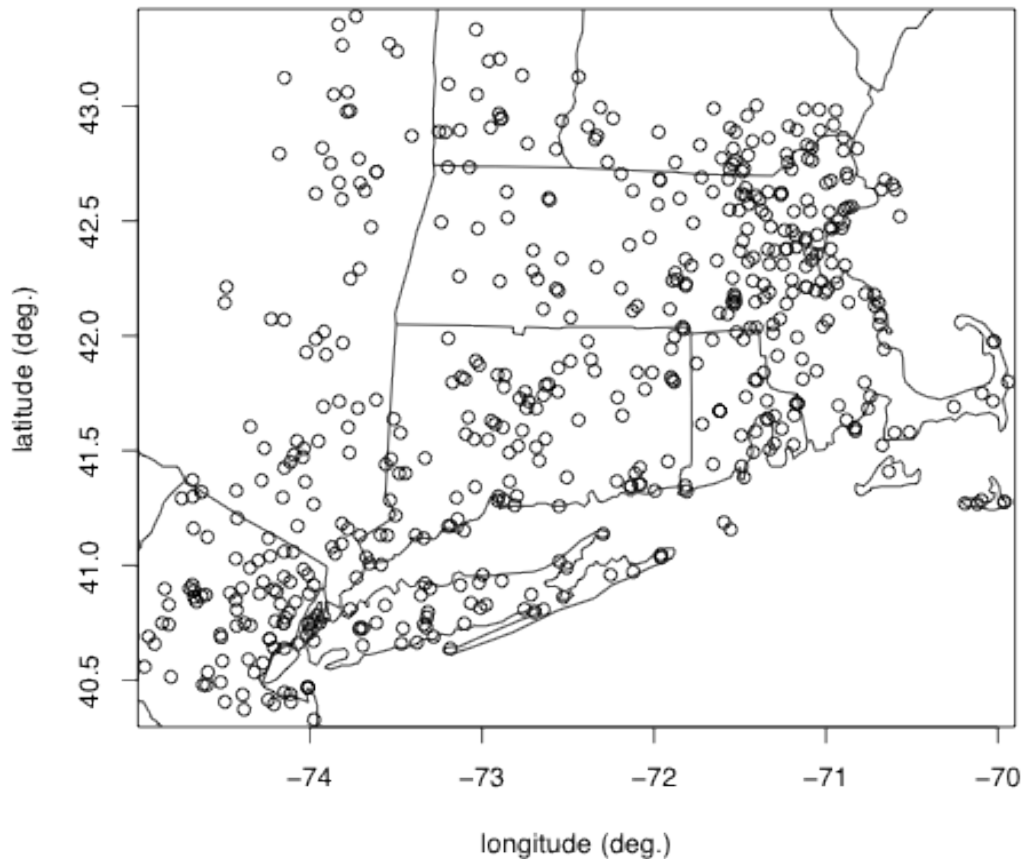


Figure 32: Locations of MADIS observations used to verify the Boston simulation.

The MADIS observations showed an overwhelming low bias that, unlike in our simulations, was exacerbated during the daytime (Figure 33). Examining just the official NWS observations, a low bias was still present, but much weaker. Observations showed much better agreement with the forecasts. Like in our study, the MADIS observations used to verify the Boston forecasts were ostensibly given at 10 m above ground level. However, verification statistics call that into question. The forecast that is present is from August 2010, but its date is unlikely to affect its relevance as our results were shown to be independent of season (Figure 12). Given the consistency between our results and the preliminary results from the Boston simulations, we are satisfied that we have identified the source of large biases as a data quality problem with the MADIS observations.




Figure 33: Domain-averaged wind speed forecasts (red) and observations (blue) from the MADIS network (left) and official NWS stations (right) for a 10-day period in August 2010 from the Boston simulations.

Comparison to results from the North Sea

The shape of the wind speed distribution from both our model simulations and the buoy observations verify well with those from the Fino 1 offshore platform in the German Bight of the North Sea (Figure 16, Figure 17, Figure 18, and Figure 34). This comparison is tangential to the task at hand and is mentioned just as a way to compare our results to other offshore observations. This serves to demonstrate that our model forecasts share a unimodal distribution shape that is expected for offshore winds. This gives us additional confidence in the ability of our simulations to produce a realistic wind climatology. While the details of the wind speed distribution, such as average wind speed and seasonal dependence, will differ between the German Bight and U.S. East Coast, we expect the general shape of the distribution to remain similar, as it should be largely dictated by the neutrally stable boundary layer.




Figure 34: Wind speed histogram from the Fino 1 platform (Beeken et al., 2008).

List of Acronyms and Abbreviations

- AFWA – U.S. Air Force Weather Agency
- ARW – Advanced Research WRF
- ASOS – Automated Surface Observing System, the official NWS weather station network
- BOEMRE – Bureau of Ocean Energy Management, Regulation, and Enforcement
- DEOS – Delaware Environmental Observing System
- KGED – Georgetown, DE surface weather station
- KOXB – Ocean City, MD surface weather station
- KWWD – Wildwood, NJ surface weather station
- MADIS – Meteorological Assimilation Data Ingest System, an NCEP product
- MET – Model Evaluation Tools from the NCAR Development Testbed Center
- MYJ – Mellor-Yamada-Janjic boundary layer scheme
- MYNN - Mellor-Yamada-Nakanishi-Niino boundary layer scheme
- NARR – North American Regional Reanalysis
- NCAR – National Center for Atmospheric Research
- NCEP – National Centers for Environmental Prediction, part of the NWS and NOAA
- NOAA – National Oceanographic and Atmospheric Administration, part of the Department of Commerce
- NWS – National Weather Service, part of NOAA
- OCS – Outer Continental Shelf
- WRF – Weather Research and Forecasting model

References

Adams-Selin, R.D., R. Johnson, and S. van den Heever, 2011: Impact of microphysics parameterizations on simulation of 13 March 2003 bow echo. *Preprints, 20th Conf. on Numerical Weather Prediction*, Seattle, WA, Amer. Meteor. Soc., 13B.5.

Courant, R., K. Friedrichs, H. Lewy, 1928: Über die partiellen Differenzgleichungen der mathematischen Physik. *Mathematische Annalen*, **100**, 32-74.

Ferrier, B. S., Y. Jin, Y. Lin, T. Black, E. Rogers, and G. DiMego, 2002: Implementation of a new grid-scale cloud and precipitation scheme in the NCEP Eta model. *Preprints, 15th Conf. on Numerical Weather Prediction*, San Antonio, TX, Amer. Meteor. Soc., 280-283.

Grell, G.A. and D. Devenyi, 2002: A generalized approach to parameterizing convection combining ensemble and data assimilation techniques, *Geoph. Res. Let.*, **29**, NO 14., 10.1029/2002GL015311.

Hong, S.Y. and Y. Noh, 2006: A new vertical diffusion package with an explicit treatment of entrainment processes. *Mon. Wea. Rev.*, **134**, 2318–2341.

Hsu, S. A., E. A. Meindl, and D. B. Gilhousen, 1994: Determining the Power-Law Wind-Profile Exponent under Near-Neutral Stability Conditions at Sea. *J. Appl. Meteor.*, **33**, 757-765.

Iacono, M. J., J. S. Delamere, E. J. Mlawer, M. W. Shephard, S. A. Clough, and W. D. Collins, 2008: Radiative forcing by long-lived greenhouse gases: Calculations with the AER radiative transfer models. *J. Geophys. Res.*, **113**, D13103.

Janjic, Z. I., 2000: Comments on "Development and Evaluation of a Convection Scheme for Use in Climate Models". *J. Atmos. Sci.*, **57**, 3686.

Janjic, Z. I., 2002: Nonsingular Implementation of the Mellor-Yamada Level 2.5 Scheme in the NCEP Meso model. *NCEP Office Note No. 437*, 61 pp.

Kain, John S., 2004: The Kain-Fritsch Convective Parameterization: An Update. *J. Appl. Meteor.*, **43**, 170–181.

Klemp, J. B., J. Dudhia, and A. D. Hassiotis, 2008: An Upper Gravity-Wave Absorbing Layer for NWP Applications. *Mon. Wea. Rev.*, **136**, 3987–4004.

Janjic, Z. I., 2003: A Nonhydrostatic Model Based on a New Approach. *Meteorology and Atmospheric Physics*, **82**, 271-285.

Klemp, J. B., J. Dudhia, A. D. Hassiotis, 2008: An Upper Gravity-Wave Absorbing Layer for NWP Applications. *Mon. Wea. Rev.*, **136**, 3987–4004.

Mesinger, F, G. DiMego, E. Kalnay, K. Mitchell, P. C. Shafran, W. Ebisuzaki, D. Jovic, J. Woollen, E. Rogers, E. H. Berbery, Y. Fan, R. Grumbine, W. Higgins, H. Li, Y. Lin, G. Manikin, D. Parrish, and W. Shi, 2006: North American Regional Reanalysis. *Bull. Amer. Meteor. Soc.*, **87**, 343–360.

Mlawer, E.J., S.J. Taubman, P.D. Brown, M.J. Iacono and S.A. Clough, 1997: RRTM, a validated correlated-k model for the longwave. *J. Geophys. Res.*, **102**, 16,663-16,682.

Nakanishi, M. and H. Niino, 2004: An improved Mellor-Yamada level-3 model with condensation physics: Its design and verification. *Bound. Layer Meteor.*, **112**, 1–31.

Nehrkorn, T. R., J. Henderson, M. Leidner, M. Ellis, A. Maher, and J. Eluszkiewicz, 2011: Modeling the urban circulation in the Salt Lake City area using the WRF urban canopy parameterization. *Preprints. Special Symposium on Applications of Air Pollution Meteorology*, Seattle, WA, Amer. Meteor. Soc., P2.838.

Otkin, J. A. and T. J. Greenwald, 2008: Comparison of WRF model-simulated and MODIS- derived cloud data. *Mon. Wea. Rev.*, **136** (6), 1957–1970.

Saha, S., et al., 2010: The NCEP Climate Forecast System Reanalysis. *Bull. Amer. Meteor. Soc.*, **19**, 1015 – 1057.

Skamarock, W. C., 2004: Evaluating mesoscale NWP models using kinetic energy, *Mon. Wea. Rev.*, **132**, 3019 – 3032.

Skamarock, W. C. and J. B. Klemp, 2008: A time-split nonhydrostatic atmospheric model for weather research and forecasting applications. *J. Comput. Phys.*, **227** (7), 3465–3485.

Smagorinsky, J., 1963: General circulation experiments with the primitive equations. *Mon. Wea. Rev.*, **91**, 99–164.

Woods, B. K., and R. B. Smith, 2011: Short wave signatures of stratospheric mountain wave breaking. *J. Atmos. Sci.*, In Press, doi: 10.1175/2010JAS3634.