**Task 2 Final Report**
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 and Atmospheric and Environmental Research, Inc.**

**March 2012**

## Table of Contents

## List of Figures

## List of Tables

## 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 potential wind patterns to quantify wind energy potential. A simulated 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 modeled climatology with 10-minute temporal and 500 m spatial resolution for an area on the OCS. Task 1 was 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 is applied to develop the full 30-year time series. The report summarizes the results of 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 ran a number of numerical weather simulations to represent typical conditions using the Weather Research and Forecasting (WRF) model (Skamarock and Klemp, 2008). A WRF model configuration that simulates the wind speed with little bias compared to buoy measurements on the OCS (compared to measurement from the one available offshore buoy) was presented in the Task 1 Final Report and is summarized in this report.

In Task 2, we take the recommended model configuration from Task 1 and use it to develop a full 30-year modeled climatology of wind resources on the OCS offshore of Delaware Bay. The simulated climatology is built by taking the historical reanalysis dataset as the background state and adding perturbations from a set of high-resolution WRF simulations. The WRF simulation to use for the perturbation at each time is chosen by matching the reanalysis dataset at each time with the most similar modeled time, as is explained later. As there are no buoys in the area of the OCS of interest, direct validation is not possible. However, this simulated

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

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 2. The methodology is presented in the next section.

## B. Methodology

The methodology described in this section is mostly a restatement of that presented in the Task 1 Final Report (MMI and AER, 2011).

### 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, 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.

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 provide the background state for the modeled 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 historical model forecast.

The reanalysis most applicable to our purposes is the North American Regional Reanalyis (NARR) from the National Centers for Environmental Prediction

(NCEP), part of the National Weather Service (NWS) and National Oceanic and Atmospheric Administration (NOAA) (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.

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 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 adjust 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 encloses the OCS area of interest (AOI) 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 is used in Task 2 to construct the full 30-year modeled climatology is shown in Figure 1.

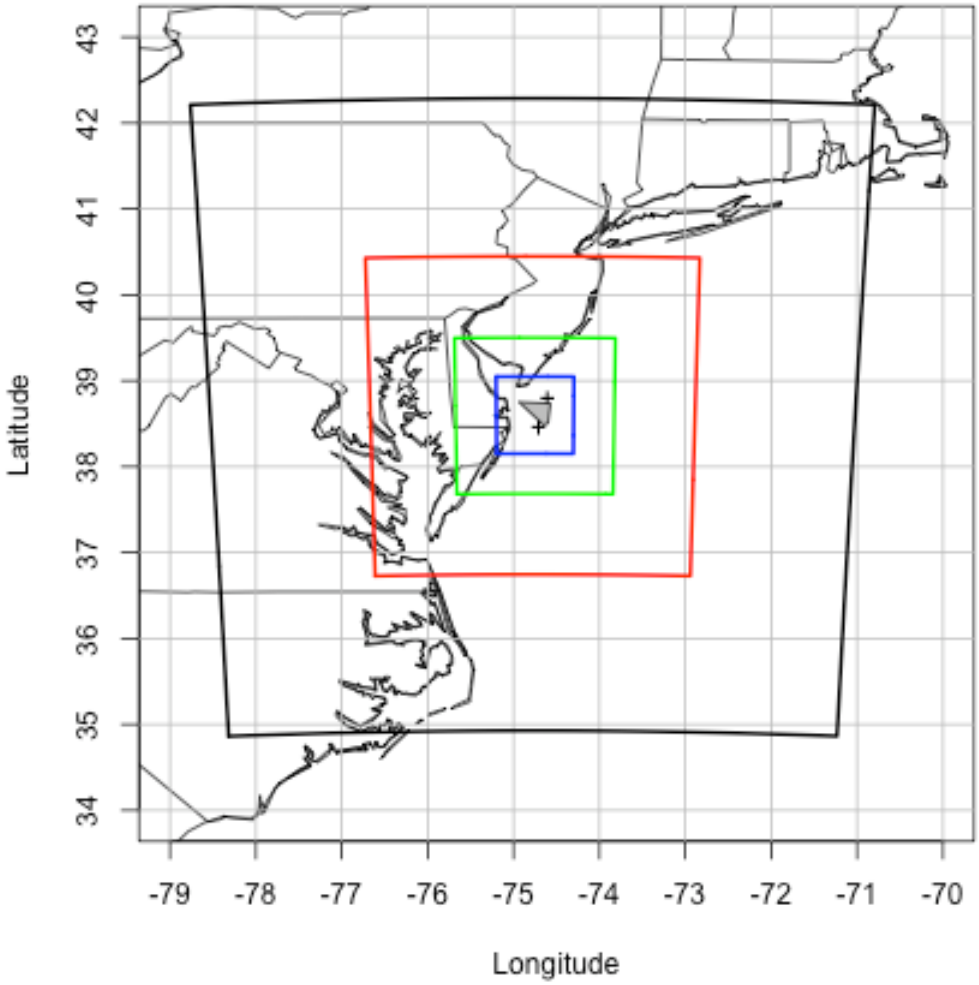


**Figure 1: Map of the four WRF domains to be used in Task 2 with grid spacing of 500 m (blue), 1.5 km (green), 4.5 km (red), and 13.5 km (black). Also shown are the area of interest (AOI; shaded), and the location of the two closest buoys (crosses).**

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 period. Both buoys are outside of the boundaries of the AOI, but are included nearby in the innermost model nest. Of those two buoys, only one was available for the sensitivity study completed in Task 1. A summary of domain sizes is given 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.**

| *Grid spacing (km)* | *Number of points West-East* | *Number of points South-North* | *Time step (seconds)* |
|---|---|---|---|
| 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. This challenge was discussed and recommendations for Task 2 were made in the Task 1 Final Report. The reader should refer to that report for further details and a brief explanation of the roles of each scheme. The options used in Task 2 are outlined below:

A. *Cloud microphysics*: This controls how water in its various phases (rain, vapor, suspended droplets, ice, etc) are treated. The choice of a microphysical scheme affects the details of cloud physics. Its influence on the wind field would be limited to areas near moist convection like thunderstorms. 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).
B. *Radiation*: 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. The radiation scheme is updated every 10 min.
C. *Boundary layer*: This parameterizations controls the transfer of momentum, heat, and moisture between the “surface layer” and “free atmosphere”. This is essentially the layer of the atmosphere that responds on an hourly scale. The influence of the boundary layer scheme was the central focus on Task 1. We use 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).
D. *Surface layer*: This specifies how the surface (ground and ocean) interacts with the atmosphere at its boundary. This scheme is dictated by our choice of the boundary layer scheme (Janjic, 2002).

E. *Land surface*: This specifies the characteristics of the land surface (e.g., available moisture, roughness, etc.). Based upon the limitations of our previous choices, we use the Rapid Update Cycle (RUC) land surface model (Smirnova et al., 2000).
F. *Cumulus parameterization*: This specifies how to handle cumulus clouds that are too small to be directly resolved by the model. This parameterization has its largest effects well above the altitudes of interest to this study. 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), the Kain – Fritsch scheme (Kain, 2004) is used to represent cumulus convection. The inner domains have high enough resolution to represent moist convective processes explicitly and in those domains no cumulus scheme is used.
G. *Mixing options*: This specifies how to control the mixing of heat, moisture, and kinetic energy between model grid cells to account for the effect of unresolved turbulence. 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.
H. *Damping option*: 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, as was shown in Task 1. The state-of-the-science damping layer of Klemp et al. (2008) is used to handle the upper boundary conditions.

## Model resolution

When considering WRF simulations, it is important to remember that for a given model resolution, not all processes are fully resolved. For example, given our 500 m grid spacing it should be clear that processes on a scale less than 500 m, like the wake of wind turbine, cannot be resolved. Processes must be much larger than the grid spacing to be accurately represented. 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 physical 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 misses 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 are not fully resolved.

Our simulations capture the large-scale sustained winds, which are defined by the World Meteorological Organization as the 10-minute average. We do not consider the limits of our resolution to be a significant limitation as temporal resolution of the final product is greater than the shortest resolvable process. Analysis of our simulations from Task 2 shows that the there is almost no difference between the 10-minute averaged wind speeds and those with 1-minute or instantaneous averaging (Figure 2, Figure 3).

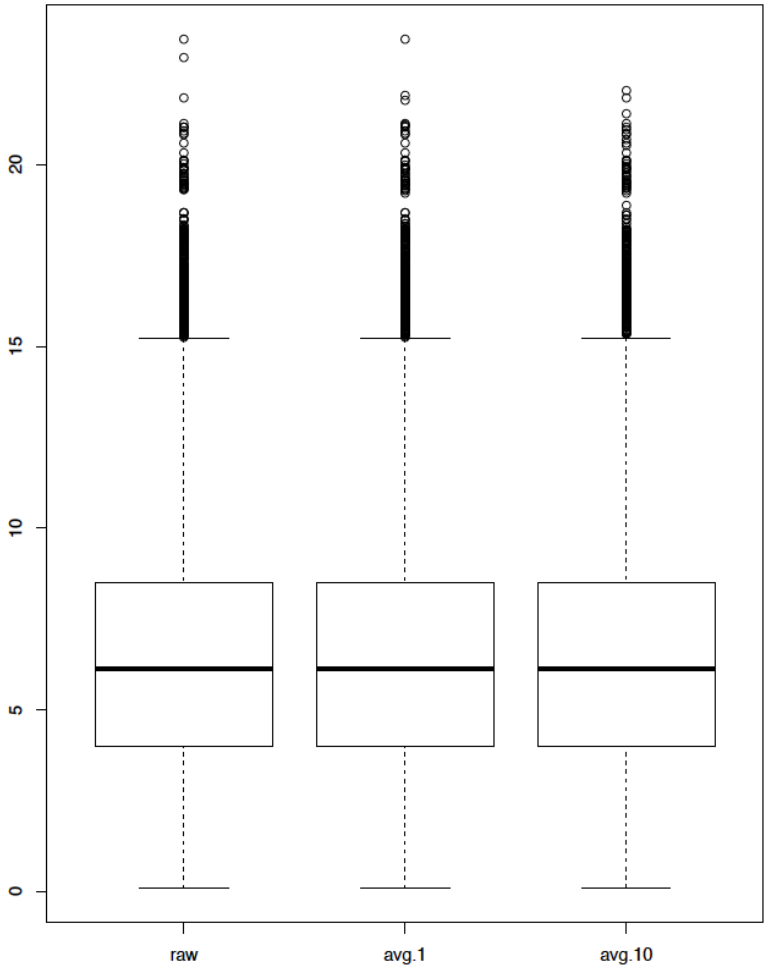


**Figure 2: Box plot of wind speed forecasts (m/s) for "instantaneous" (raw), 1 min, and 10 min averaging. Note that all three distributions are nearly identical. The thick black lines are the median values (50% quantile). The box outline represents the 25-75% quantiles. The outer whiskers encompass the range of data points, up to a distance of 1.5 times the interquartile range from the box (for a normal distribution, this includes 99% of all data points). Outliers (data points beyond the whiskers) are shown with circles.**

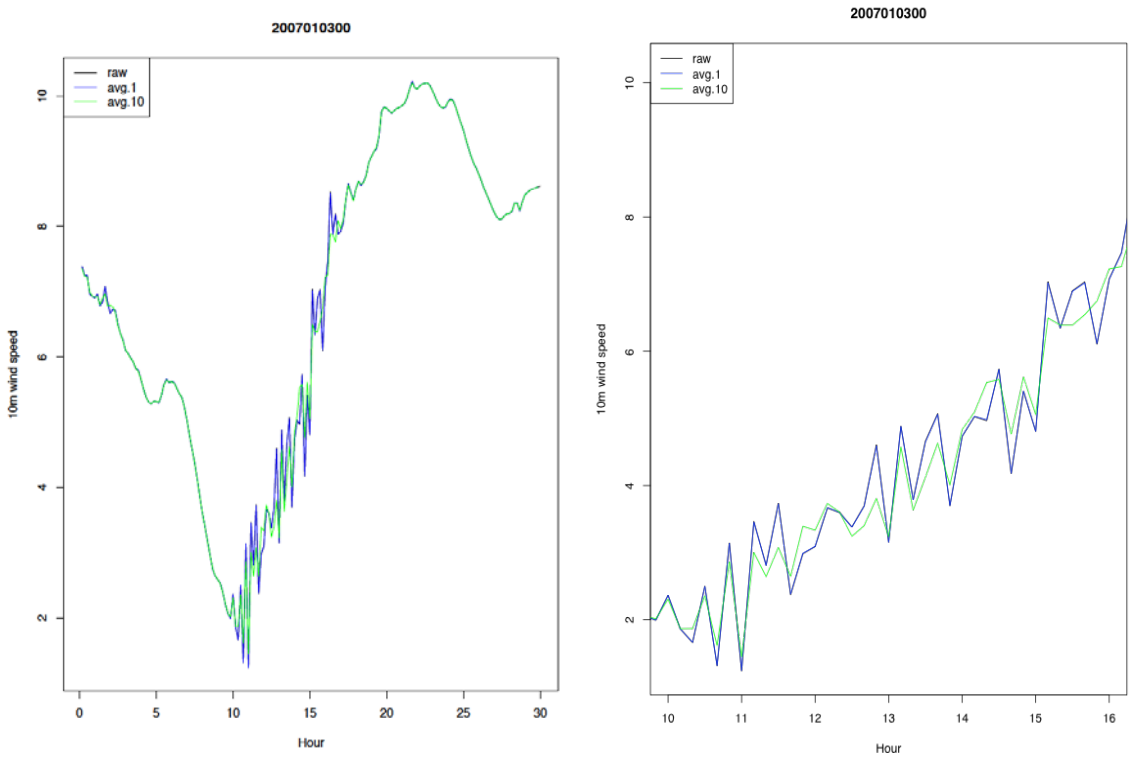


**Figure 3: Example time series of a 30-hour wind speed forecast (m/s) for instantaneous (raw), 1 min, and 10 min averaging. The right panel is a detailed view of the part of the forecast with noticeable differences between the 1 and 10 min winds. Note that the 1 min and instantaneous WRF winds are nearly identical, and that while there are some differences between 1 min and 10 min winds at some times during the day, there are no systematic biases (consistently lower or higher wind speeds), consistent with the overall statistics shown in Figure 2.**

We have focused our simulations on the altitudes where wind generation occurs in the first few hundred 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".

For our purposes, that means we try to make the model level spacing match the spacing needed in our final product, even if the levels themselves are different. We define a "typical day" as a quiet summer day. This is the time when we expect the model levels to be spaced furthest apart due to thermal expansion of the air column. During the winter when the air is colder, it is also denser so that our mass-based vertical coordinates will be located closer together.

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 troposphere[2], 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.

### Choice of testing period

In order to most faithfully represent the actual variability expected in a full climatology, we sought days for our WRF simulations that were independent samples (i.e., far enough apart to not be correlated) and not preselected to meet some preconception. Within the scope of the available computing resources, we are limited to simulating three years with WRF. Working under the assumption that more recent data will be of higher quality, we chose to sample from calendar years 2007-2009 for our WRF simulations in Task 2. Compared to the year 2010, which was used in Task 1 (Figure 5), the years 2007-2009 (Figure 4) show relatively more frequent and stronger northeasterly winds. Having a more diverse distribution in the years 2007-2009 will aid our sampling methodology by providing a wider variety of WRF simulations to use to construct the model time series.

Wind roses for the NARR 10 m winds over the entire 3-year period and the just the subsample simulated with WRF are shown in Figure 4 and Figure 6 respectively. To construct the wind roses, the NARR analyses are interpolated to the 500 m grid and the central grid point is used. The wind roses can be thought of as a type of histogram of the 3-hour wind showing the relative frequency of a number of speed and direction bins.

The subsampled interval modeled with WRF (Figure 6) 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 strong southeasterly winds is typical of mid-latitudes. Note that all winds are named based on direction from which they come (*i.e.*, a sea breeze blows from sea toward land and a northwesterly wind blows from the northwest toward the southeast). The occurrence (or lack) of "extreme" events in the

[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.

simulated period is less important than faithfully representing the wind distribution. The sampling strategy addressed in the next section will account for the relative frequency of occurrence of specific weather conditions.

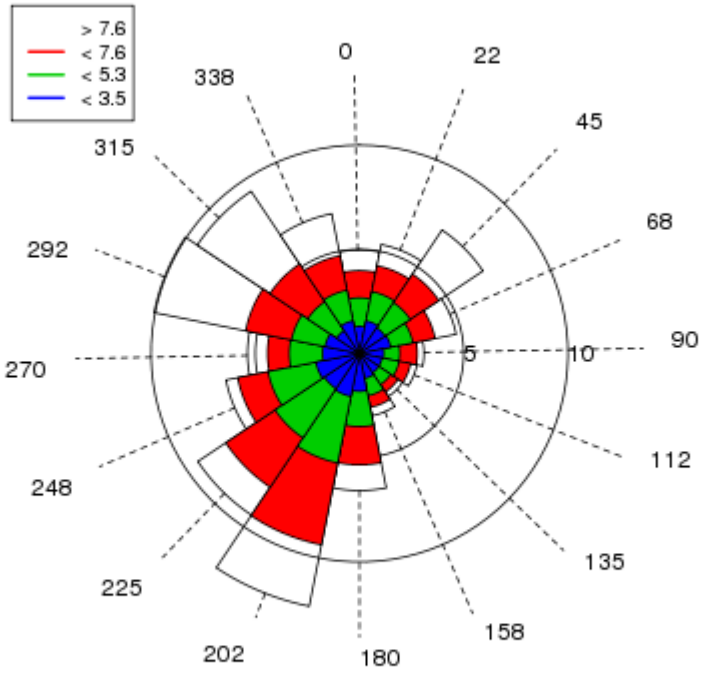


**Figure 4: Wind rose (10 m) for the AOI (shown in Figure 1) from the NARR historical reanalysis for the full calendar years 2007-2009. Each slice is representative of a directional bin (labeled in degrees from north). The length of the slice is representative of the relative frequency (%), labeled along the x-axis. The colored portions correspond to a speed (m/s) bin for each direction (i.e., green is 3.5 – 5.3 m/s).**

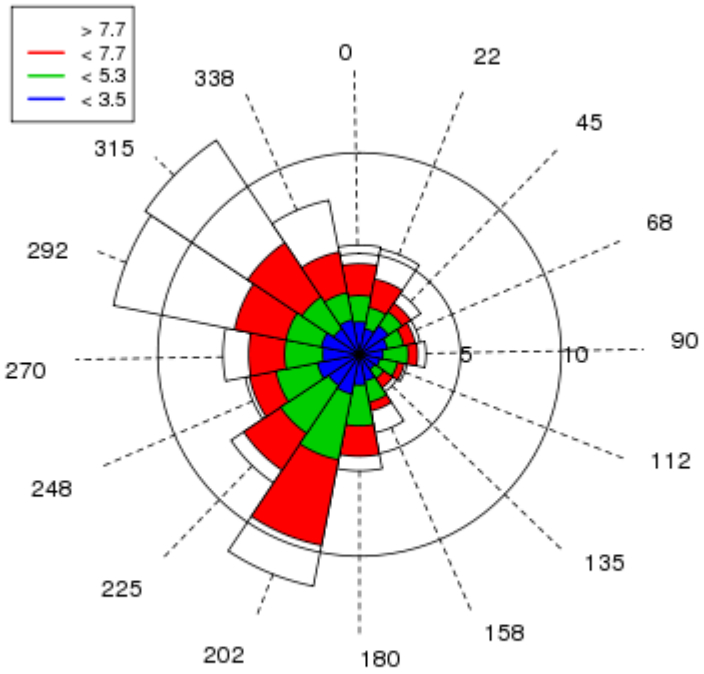


**Figure 5: Wind rose (10 m) for the AOI (shown in Figure 1) from the NARR historical reanalysis for the full calendar year 2010 which was sampled in Task 1. Each slice is representative of a directional bin (labeled in degrees from north). The length of the slice is representative of the relative frequency (%), labeled along the x-axis. The colored portions correspond to a speed (m/s) bin for each direction (i.e., green is 3.5 – 5.3 m/s).**

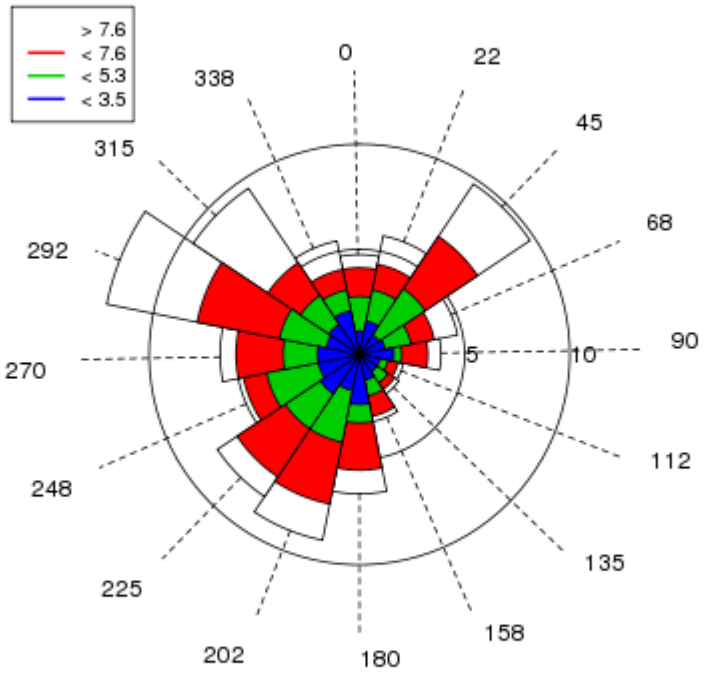


**Figure 6: Wind rose (10 m) for AOI (shown in Figure 1) from the NARR historical reanalysis for the dates used in our Task 2 simulations. Each slice is representative of a directional bin (labeled in degrees from north). The length of the slice is representative of the relative frequency (%), labeled along the x-axis. The colored portions correspond to a speed (m/s) bin for each direction (i.e., green is 3.5 – 5.3 m/s).**

## Building the full 30-year time series

Before combining the WRF simulations and NARR, the analyses must first be converted to the same underlying grid. As the two datasets are natively given at different spatial and temporal resolutions, this will involve a preprocessing step to transform the data to a common framework.

The WRF and NARR data are first reprojected and interpolated to the final grid. Both datasets are originally delivered on Lambert Conformal grids. The reprojection coordinates in each case are calculated using the PROJ.4 package, originally produced by the United States Geological Survey. The reprojections are applied using cubic spline interpolation in the horizontal and linear interpolation in the vertical.

For each 3-hour period in the full 30-year NARR, the "most similar" modeled NARR period is found by comparing the values from the central grid point in the

reprojected data. The "similarity" of two times is optimized by minimizing a mismatch parameter that we have defined. This dimensionless parameter is calculated by summing the differences in the following fields normalized by their respective standard deviations:

- Wind shear $\left(d\dot{U}/dz\right)$ (*i.e.*, a change in the wind speed or direction with height) between z = 10 – 210 m
- Bulk stability $N^2 = \frac{g}{T}\frac{dT}{dz}$ between z = 2 – 210 m
- Sea surface temperature
- Wind speed at hub height (z = 90 m)
- Wind direction at hub height (z = 90 m)

where U is the wind speed vector, N is the buoyancy frequency, g is the acceleration of gravity, and T is the air temperature. The standard deviations used to normalize the differences are taken across the full 30 years. By normalizing each of the five terms by its standard deviation, each term, on average across the full 30 years, contributes equal weight to the matching criterion. The wind shear and bulk stability are included in the matching criterion. Both are important as the ratio of the two, defined as the Richardson number $Ri = \frac{(dU/dz)^2}{N^2}$, controls whether or not the flow will be turbulent, i.e. turbulent flow is expected for $Ri < 1/4$.

Each simulated period is sampled at least once in the 30 years. Each simulated period has to match itself as the difference between a value and itself is by definition zero (*i.e.,* $x - x = 0$). Simulated periods that represent common weather patterns are sampled repeatedly. Figure 7 shows the number of times that each NARR period that is simulated with WRF is sampled when constructing the 30-year time series.

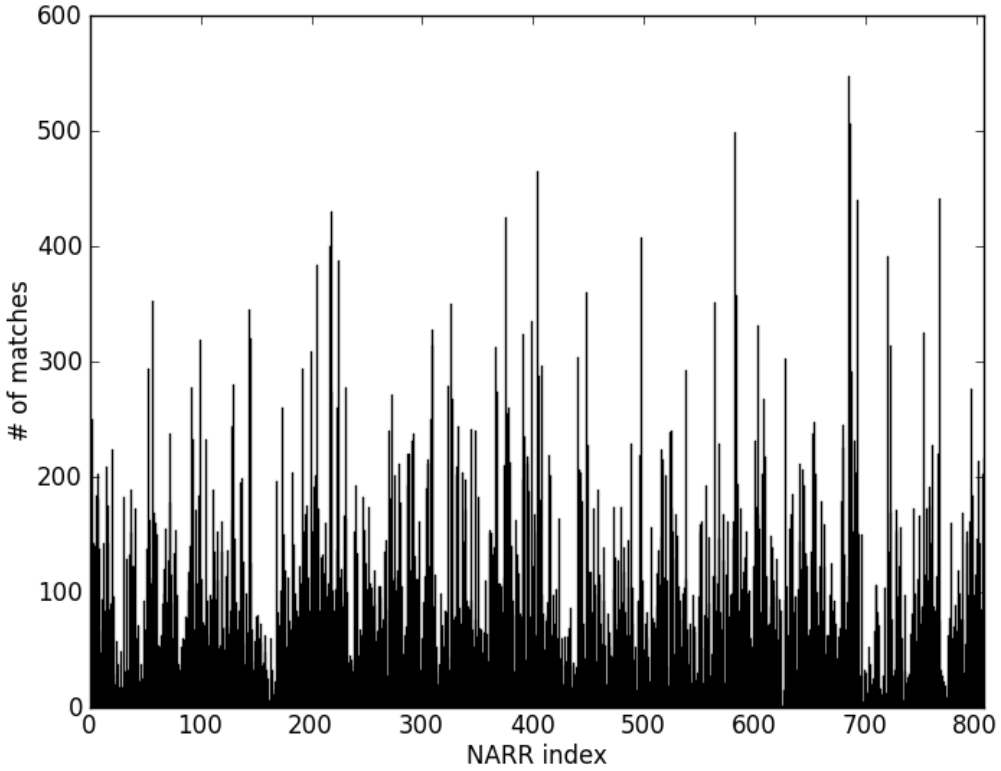


**Figure 7: Number of matches for each 3-hour NARR period simulated with WRF. Note that each period is sampled at least once, and some much more often.**

For each WRF simulation, the results are converted to perturbations from the NARR by calculating the difference between the WRF simulation results and the reprojected and interpolated (in both time and space) NARR that was used as a boundary condition for the simulation. These WRF perturbations are repeatedly sampled to construct the 30-year time series. The background atmosphere for each 3-hour interval in the 30-years is taken from the reprojected NARR. Upon that background, the WRF perturbations are added from the matching period.

## Processing hardware and time

The WRF simulations for Task 2 were conducted on an SGI Altix ICE cluster with 512 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. Running all of the simulations for Task 2 took approximately three (3) calendar days.

Post-processing of the simulations and NARR to create the 30-year time series was performed on a number of servers, including a server-grade Mac Pro with 12-processor cores and several virtual cloud-computing servers with 4 cores each. The post-processing of the NARR data to subset and interpolate the full 30-year time series to the 500 m grid and vertical levels in the AOI was completed in about one calendar day. Combining the NARR background and WRF perturbations to generate the final 30-year time series took approximately 42 hours on one core per processed month.

## C. Simulation Results

### Combined time series

After combining the WRF and NARR simulations we have a complete time series with values at each 10-minute time step in the 30-year period. As previously mentioned, the NARR prescribes the large-scale background reanalysis (a combination of regional model predictions and observations) and the WRF adds fine-scale modeled details down to the 500 m grid spacing. Figure 8 shows the wind speed at z = 90 m for an arbitrary example time from the final 30-year time series. The corresponding plot for mean annual wind speeds at z = 90 m, averaged over all years in the dataset, is shown in Figure 9 . The wind speeds in this figure vary much more smoothly in the horizontal. They are in the range of 8.6 – 9.2 m/s over the entire 500m grid, with values at or above 8.9 m/s within the AOI. The annual mean wind speed maps published by the National Renewable Energy Laboratory (Schwartz et al. 2010), which use a 0.5 m/s resolution for wind speeds above 6.0 m/s, are consistent with the values shown in Figure 9.

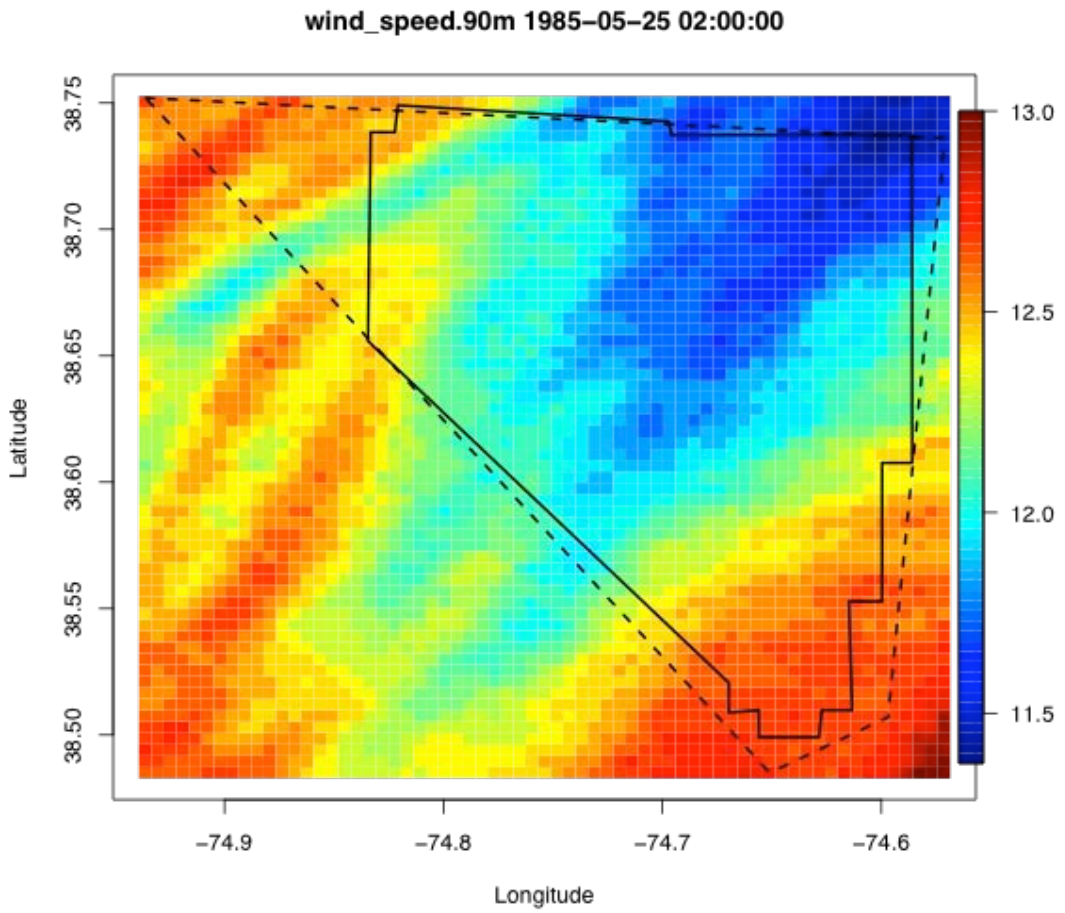


**Figure 8: Wind speed at z = 90 m from the combined forecast (NARR+WRF) at an arbitrary example time. Each pixel is a 500 m grid point. The AOI is outlined in black (initial designation as dashed, final as solid).**

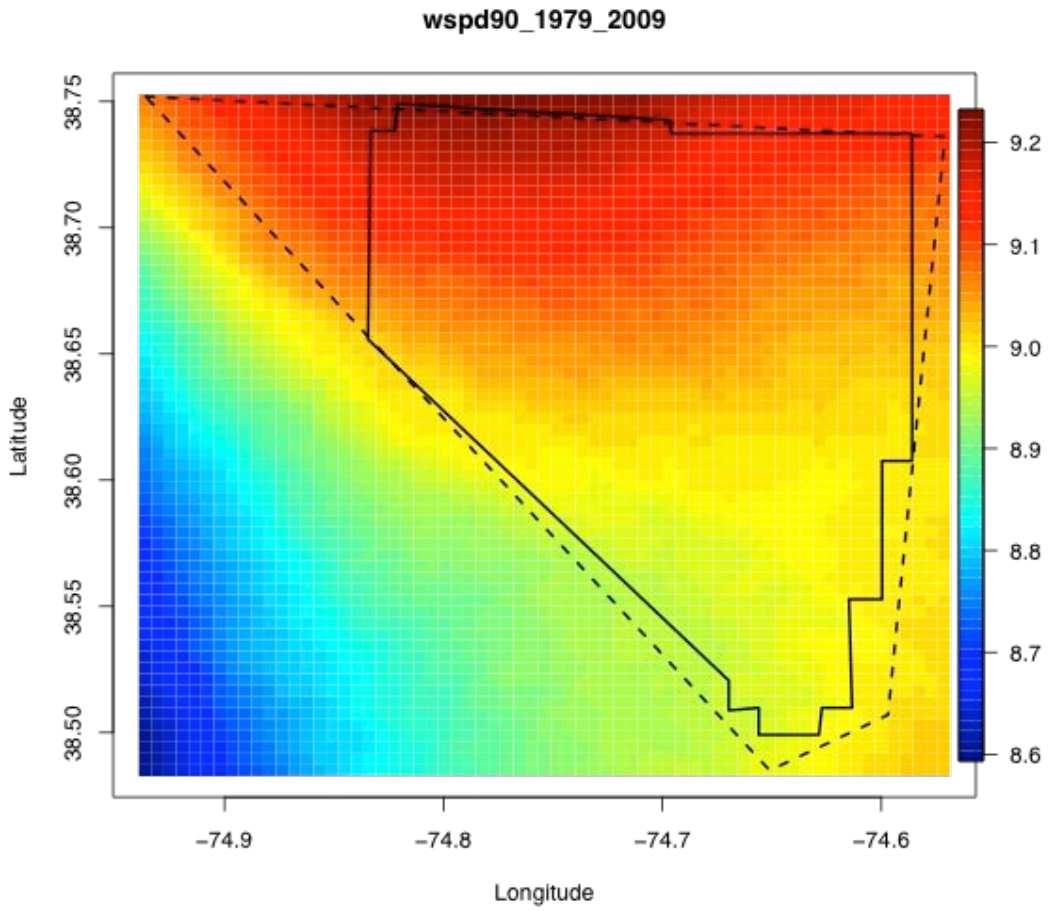


**Figure 9: Annual mean wind speed at z = 90 m from the combined forecast (NARR+WRF), averaged over all years in the simulated climatology. Each pixel is a 500 m grid point. The AOI is outlined in black (initial designation as dashed, final as solid).**

The final time series shows that winds at hub height (henceforth assumed to be z = 90 m and measured only at the central grid point) are usually in the range needed for wind energy generation. Winds are very rarely calm and also rarely exceed 25 m/s. Figure 10 shows a time series of wind speed exceedance probabilities at hub height with monthly temporal resolution for the full 30 years. The strong seasonal variation seen in the time series of Figure 10 is also evident in the monthly averages across the entire period (Figure 11). The strongest winds are clearly evident in the winter with the weakest winds in the summer. Any diurnal cycle is much more subtle (Figure 12). Considering that the ocean surface does not respond strongly to diurnal forcing this is not a surprising result.

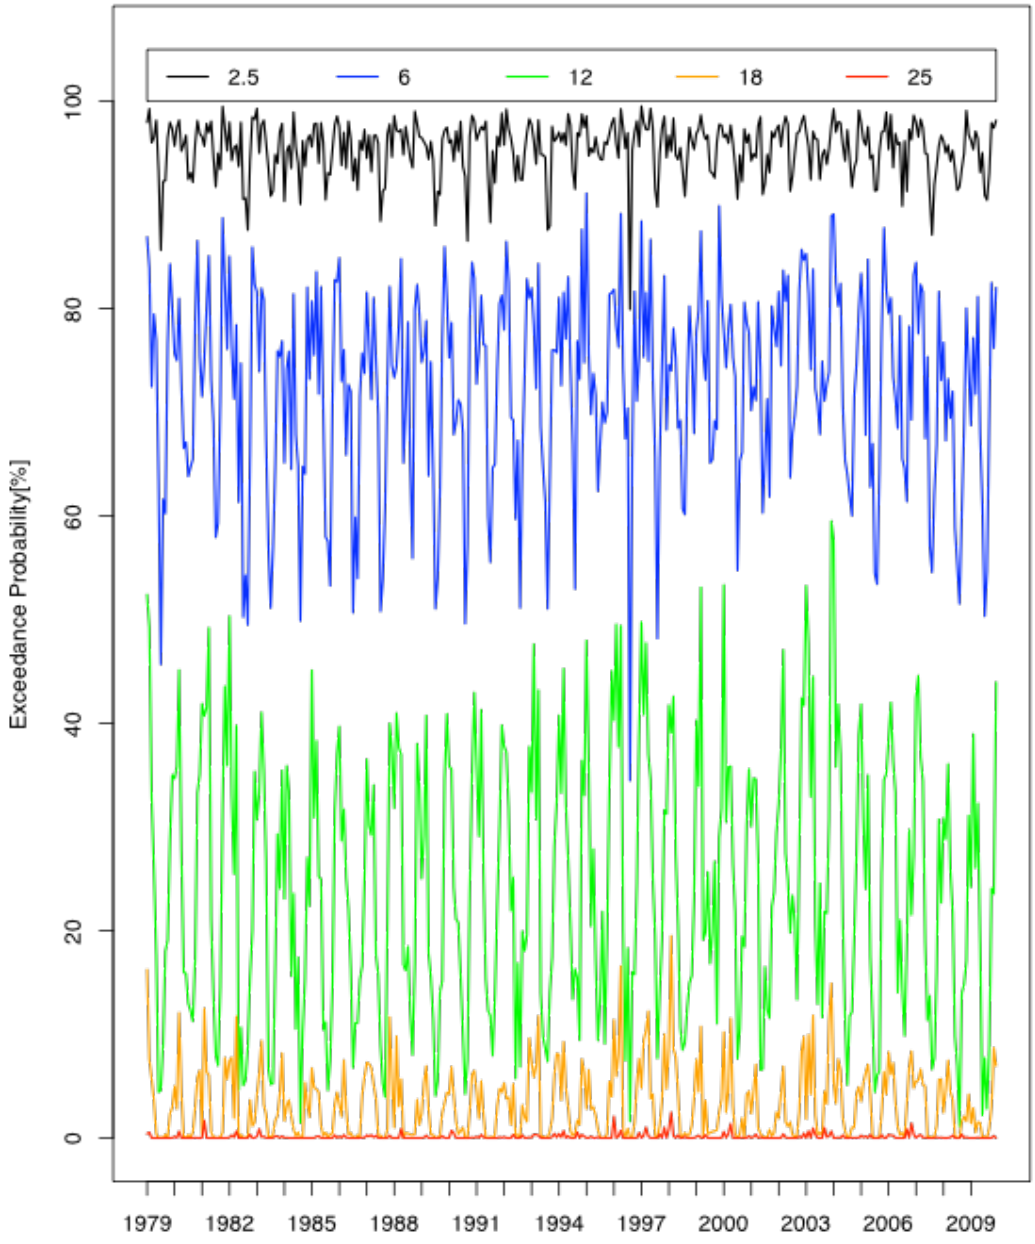


**Figure 10: Monthly exceedance probabilities for wind speed (m/s) at hub height (90 m) for the full 30 years.**

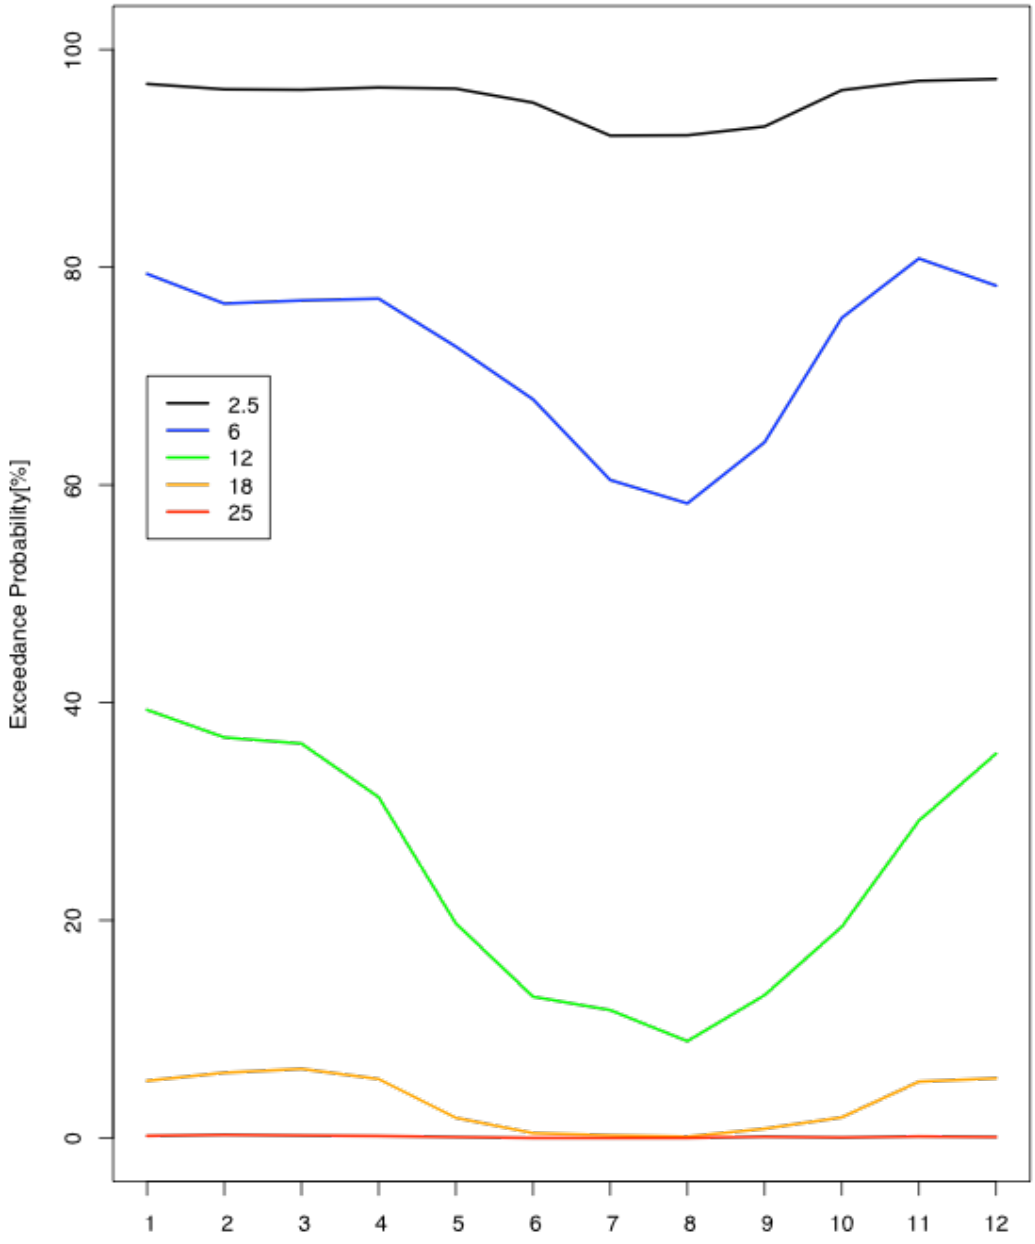


**Figure 11: Exceedance probabilities for hub height (90 m) wind speed (m/s) by month average across the full time series. Note the stronger winds in winter and weaker winds in summer.**

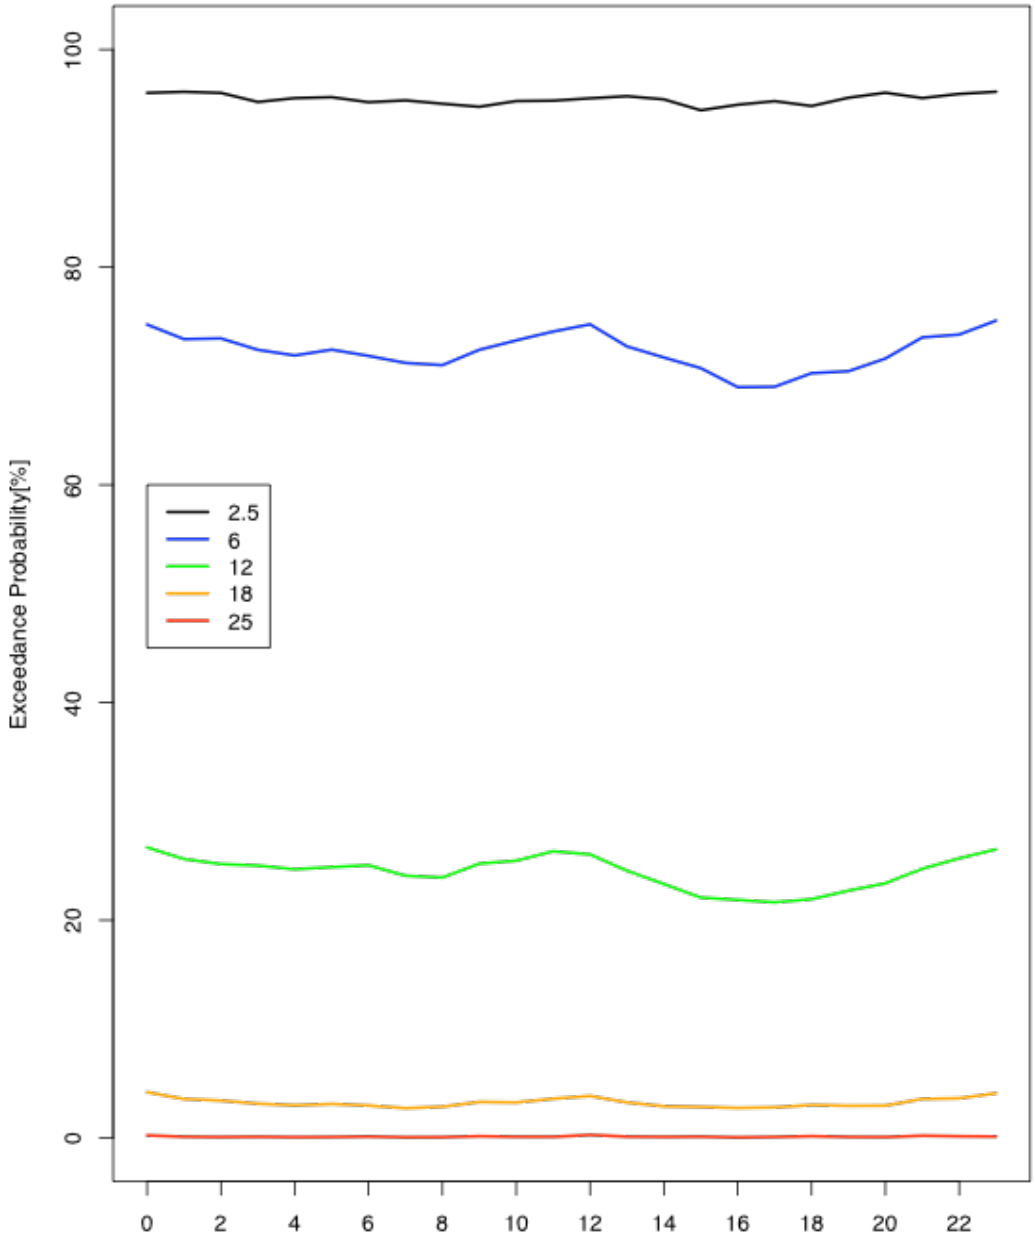


**Figure 12: Exceedance probabilities of wind speed (m/s) at hub height (m) by hour of the day (UTC / GMT / Z).**

To illustrate how to use the final data set to estimate the percentage of time a potential wind farm would operate, we again took the final wind speed values at z = 90 m for the central grid point to be representative of typical operating conditions. We defined operating conditions as hub height wind speeds in the range 2.5 – 25 m/s and 'optimal' wind speeds as 6 - 18 m/s. For these nominal values, a potential

wind farm would be capable of operating approximately 95% of the time, with optimal operations approximately 70% of the time (Figure 13). There is a clear seasonal cycle indicating a drop in operating frequency to roughly 60% during the summer (Figure 14). Once again daily cycles are much more subtle (Figure 15).

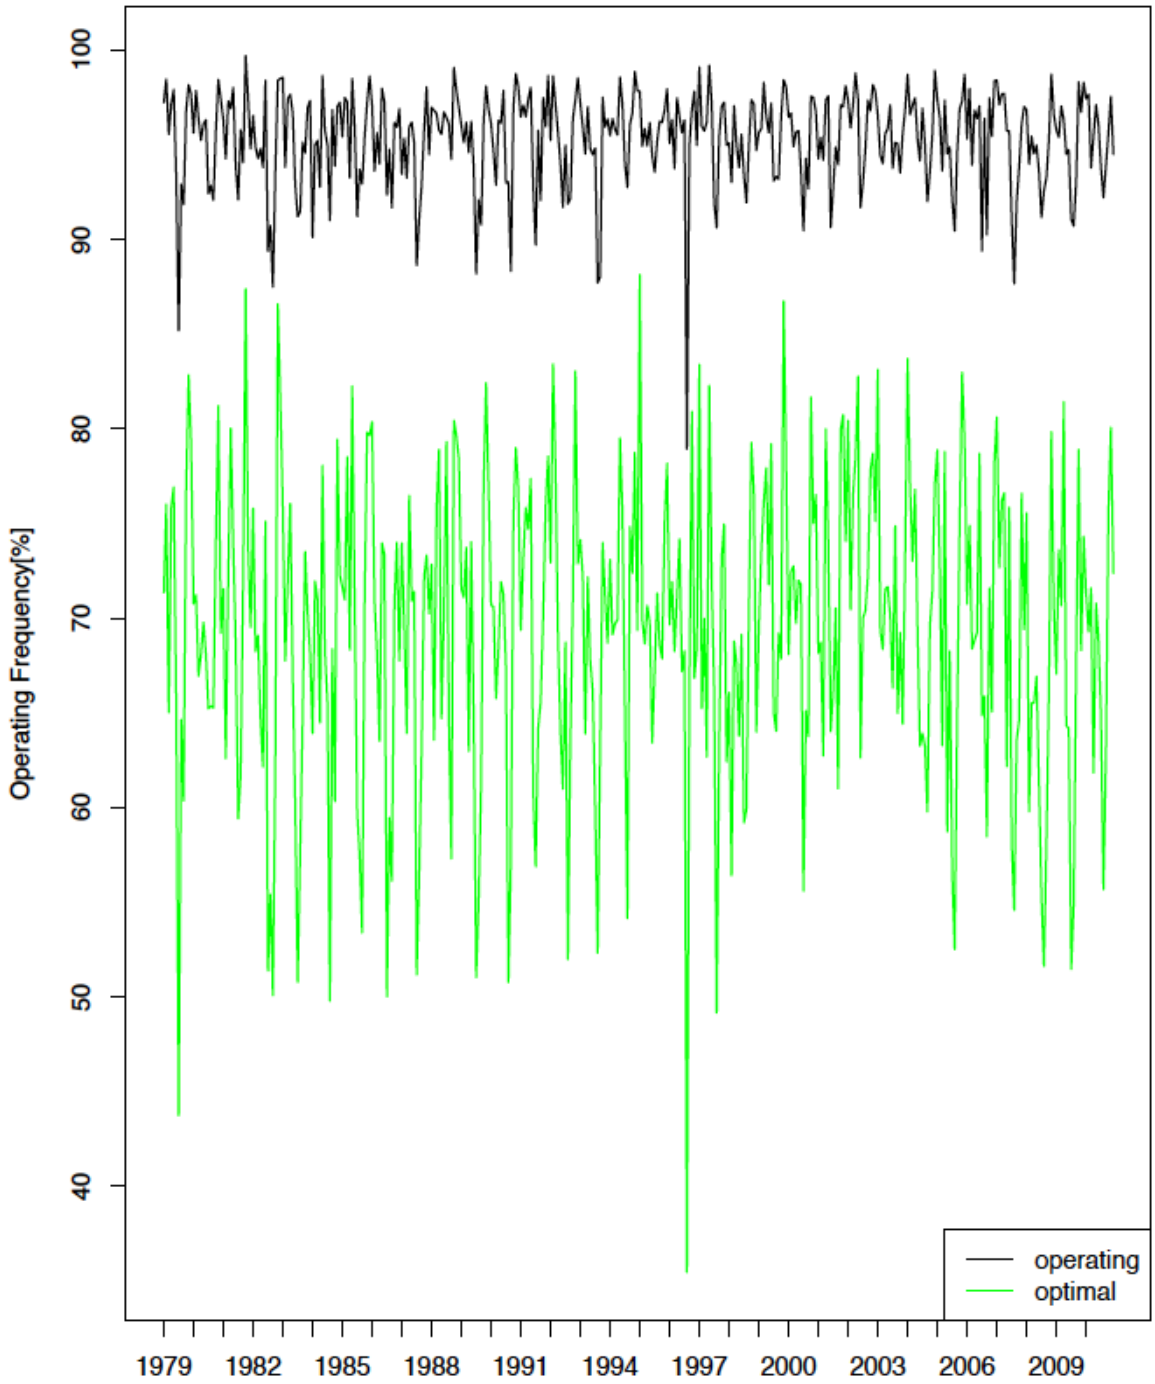


**Figure 13: Frequency of hub height wind speed in the operating (2.5 – 25 m/s) and 'optimal' (6 – 18 m/s) range. Each point in the time series is one month.**

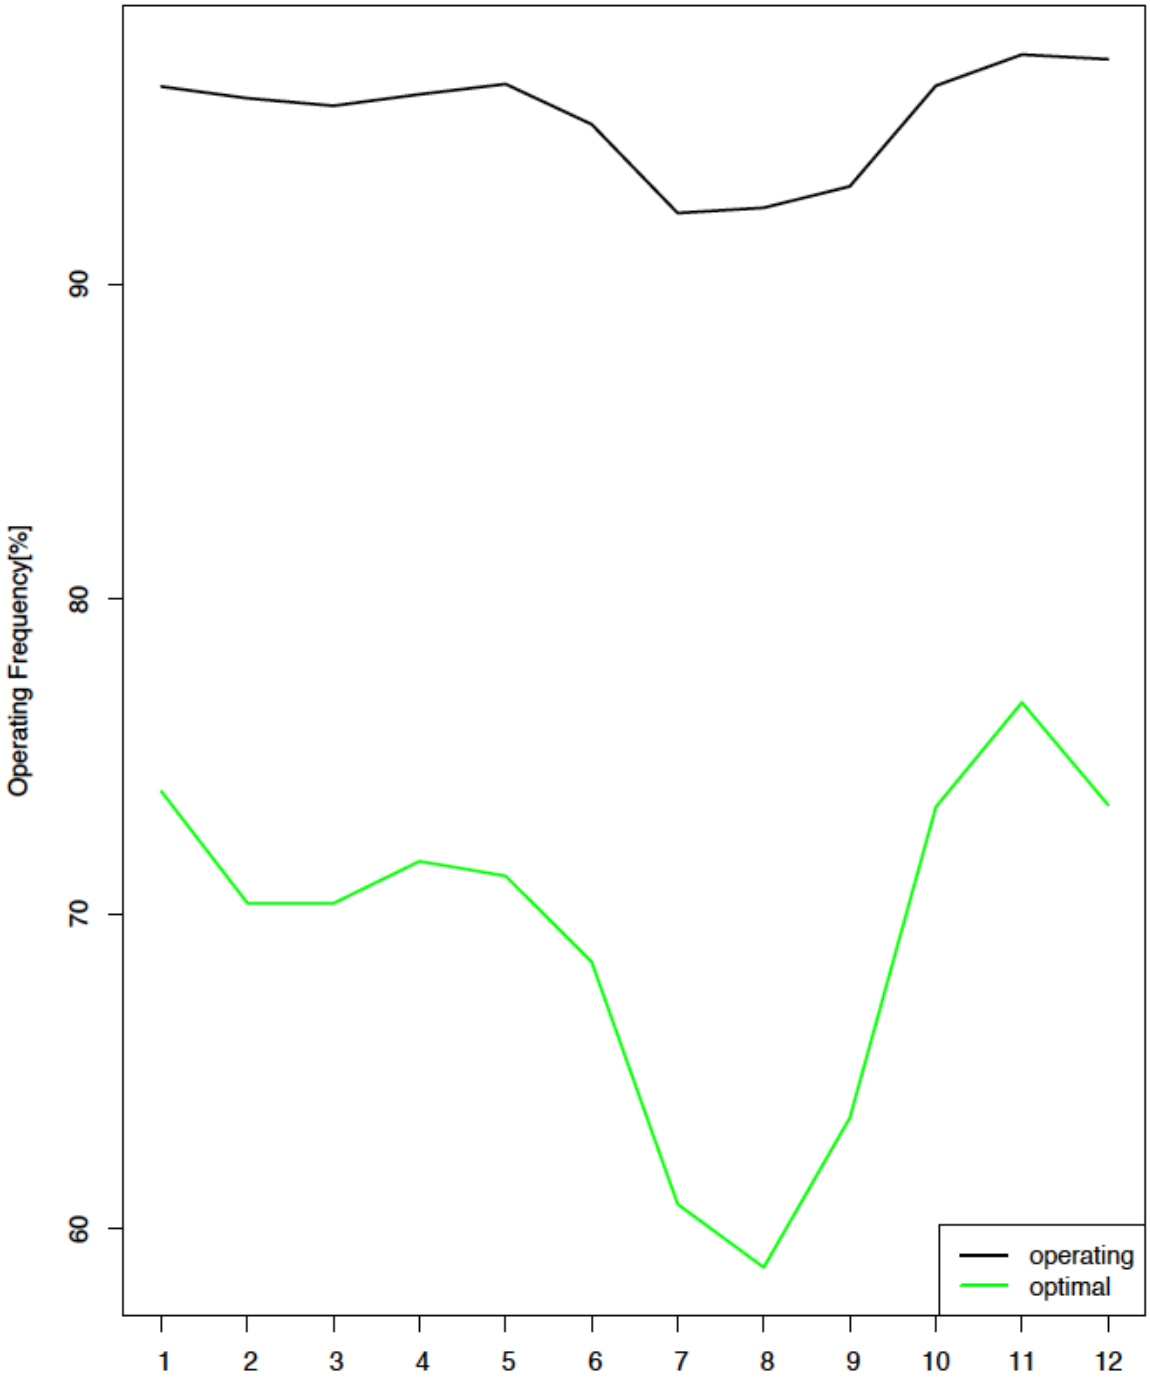


**Figure 14: Frequency of hub height wind speed in the operating (2.5 – 25 m/s) and 'optimal' (6 – 18 m/s) range for each month averaged over the full time series.**

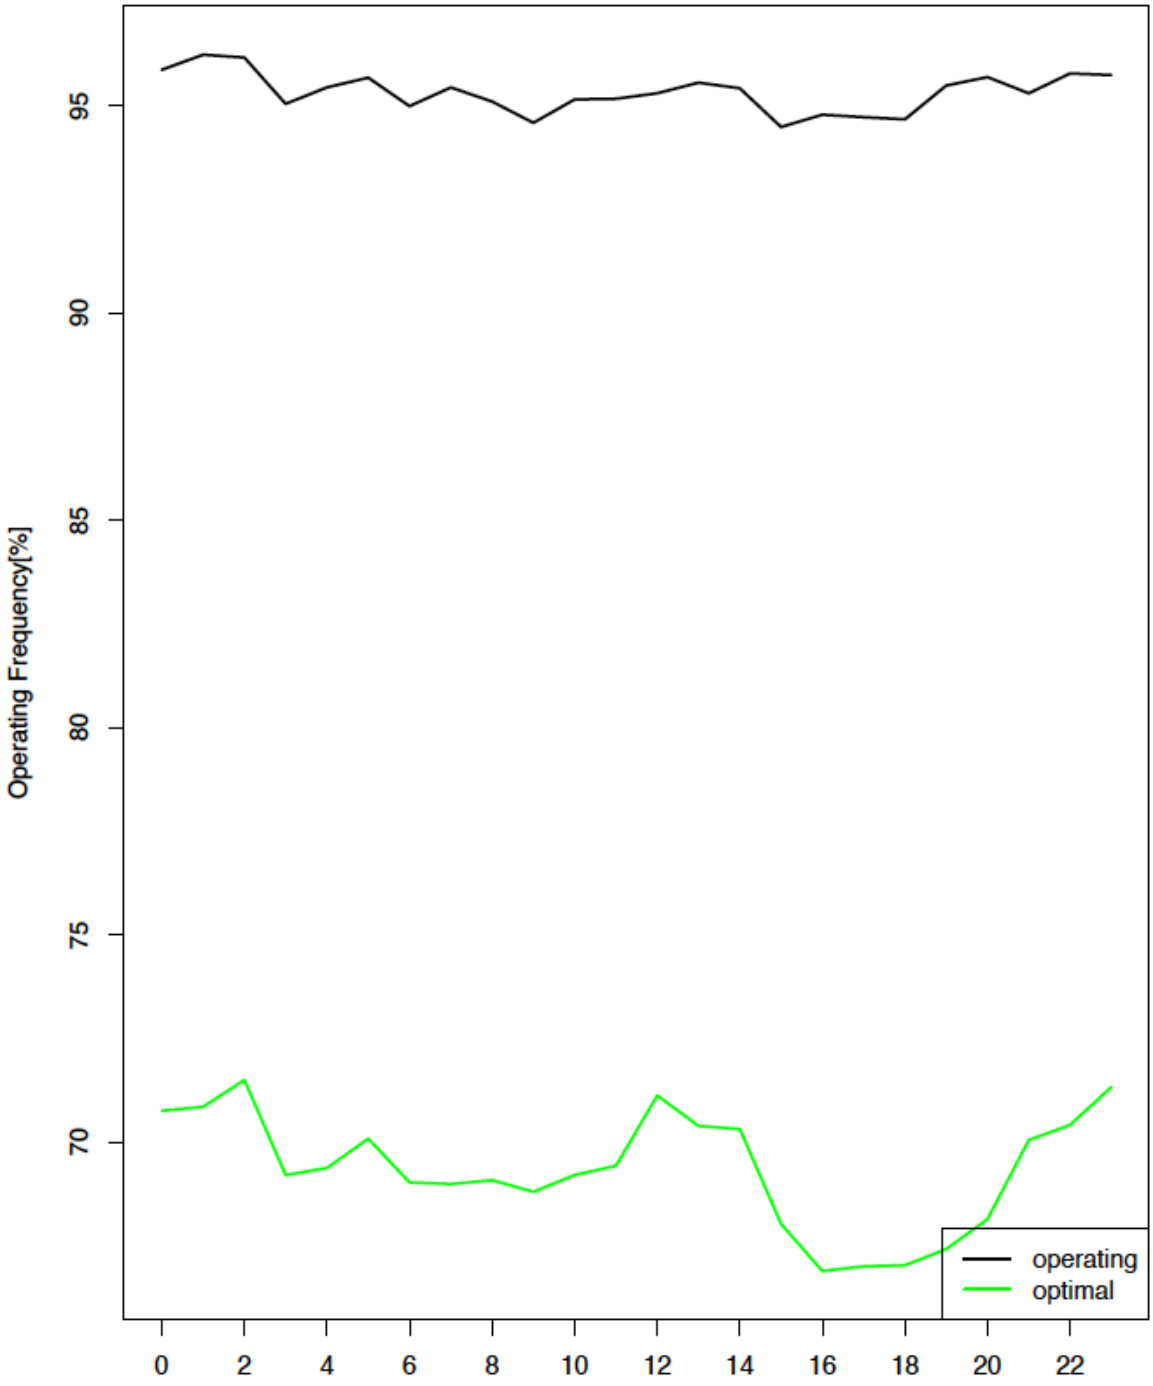


**Figure 15: Frequency of hub height wind speed in the operating (2.5 – 25 m/s) and 'optimal' (6 – 18 m/s) range for each hour of the day averaged over the full time series.**

## Validation

As there are no fixed-point observations such as buoys available in the AOI, we cannot directly verify the final time series. However, there is data for two buoys located nearby the AOI (Figure 1). Of the two buoys, only one was used for model validation and sensitivity testing in Task 1. Direct comparisons between the forecasts and observations, such as root mean square error, will have little value as we should not expect the two values to be equal. We do, however, expect the buoys and final dataset to have similar distributions in the aggregate. Comparing the wind speed distributions from the two buoys to those from the final time series, it is clear that we have a slight high bias at z = 10 m if the buoys are taken as truth. The two buoys agree well with each other, but the final time series is on average about 1 m/s higher than both buoys (Figure 16, Figure 17). For these comparisons, only time periods available at both buoy locations were considered to ensure a homogeneous sample. This limited the comparison to 3,804 hourly wind speed values between June 11, 1984 and October 28, 1992.

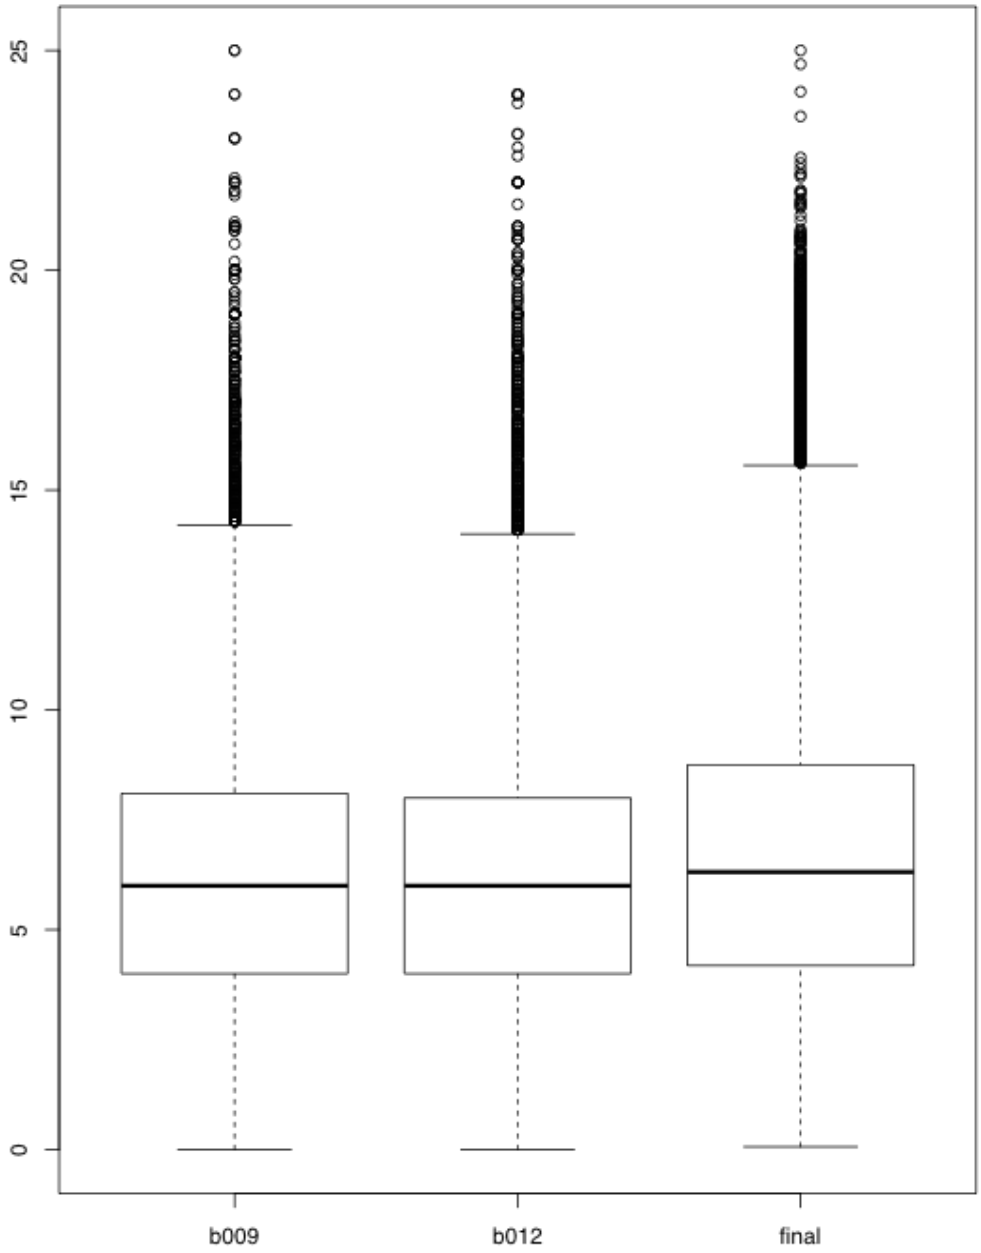


**Figure 16: Box plot of the wind speed (m/s) distributions at the two buoys (b009 and b012) and the central point of the AOI (final) time series. Note the similar distributions of the two buoys and the high bias evident in the final time series. See Figure 2 for an explanation of the box plot.**

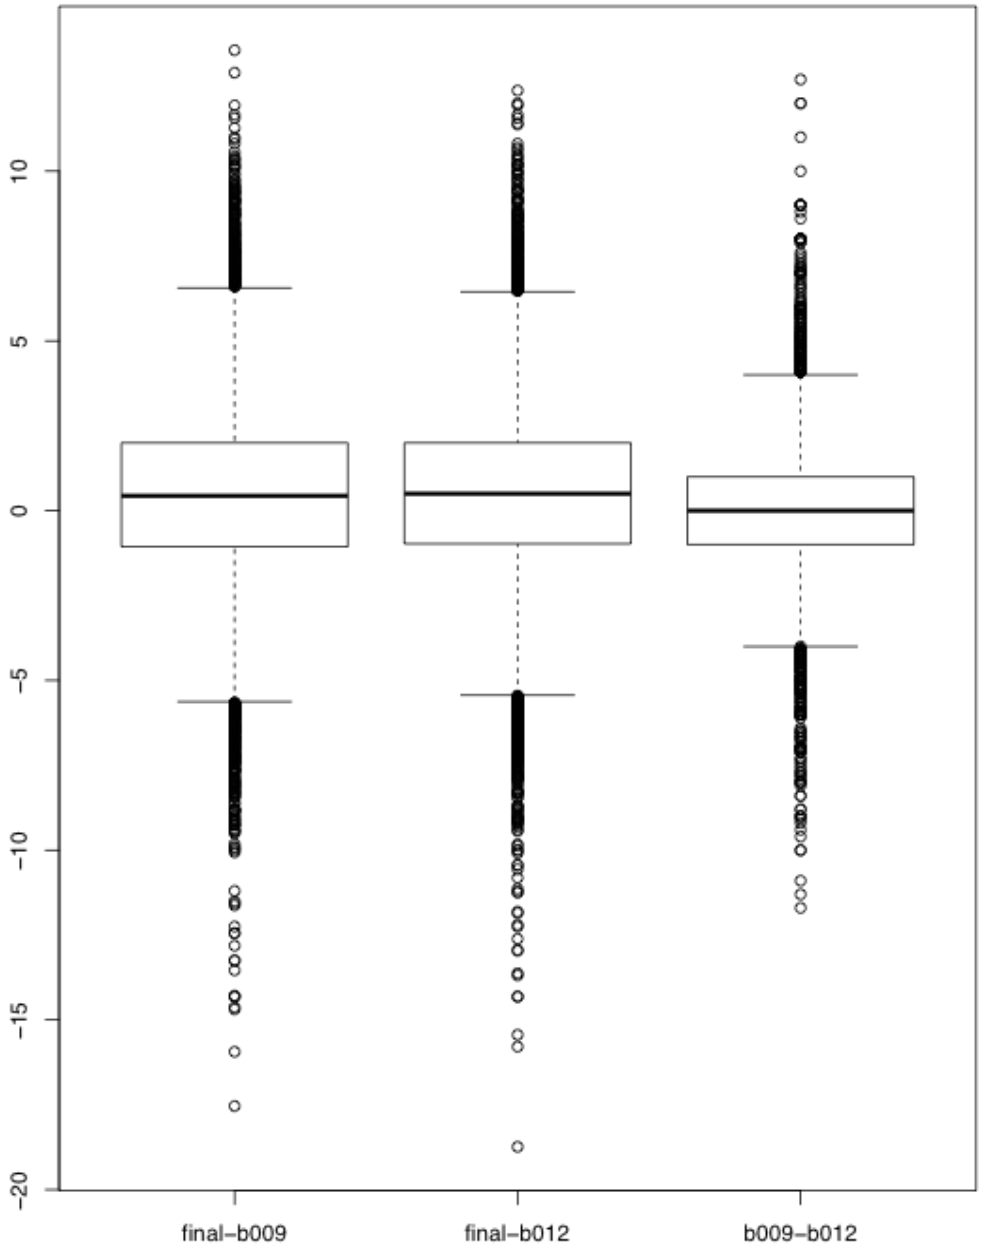


**Figure 17: Box plots for the wind speed (m/s) differences between the two buoys (buoy – final dataset) in the left two distributions. The right distribution is the difference between the two buoys. Note that the final dataset appears to have about a 1 m/s high bias compared to the buoy observations, which are consistent with each other. See Figure 2 for an explanation of the box plot.**

While the 1 m/s difference between the final dataset and the buoy observations is robust, we are not dealing with a fair comparison between the two. The two buoy observations are taken at 5 m above the sea surface whereas the model surface winds are given at the standard height of 10 m. This difference was discussed in the Task 1 Final Report where we documented that, using the standard correction recommended by the National Data Buoy Center, we expect the winds at

10 m to be on average 0.55 m/s stronger. This brings the median bias shown in Figure 17 to within the buoys' calibration error of 1 m/s. However, the wind roses from the buoys (Figure 18) do show some differences from those derived from central grid point of the AOI, but differences are not larger than those between the two buoys. There are also some noticeable differences from windroses shown in (Figure 4, Figure 6) for the NARR, as well, but these do not cover the time period, so some additional differences due to sampling can be expected.

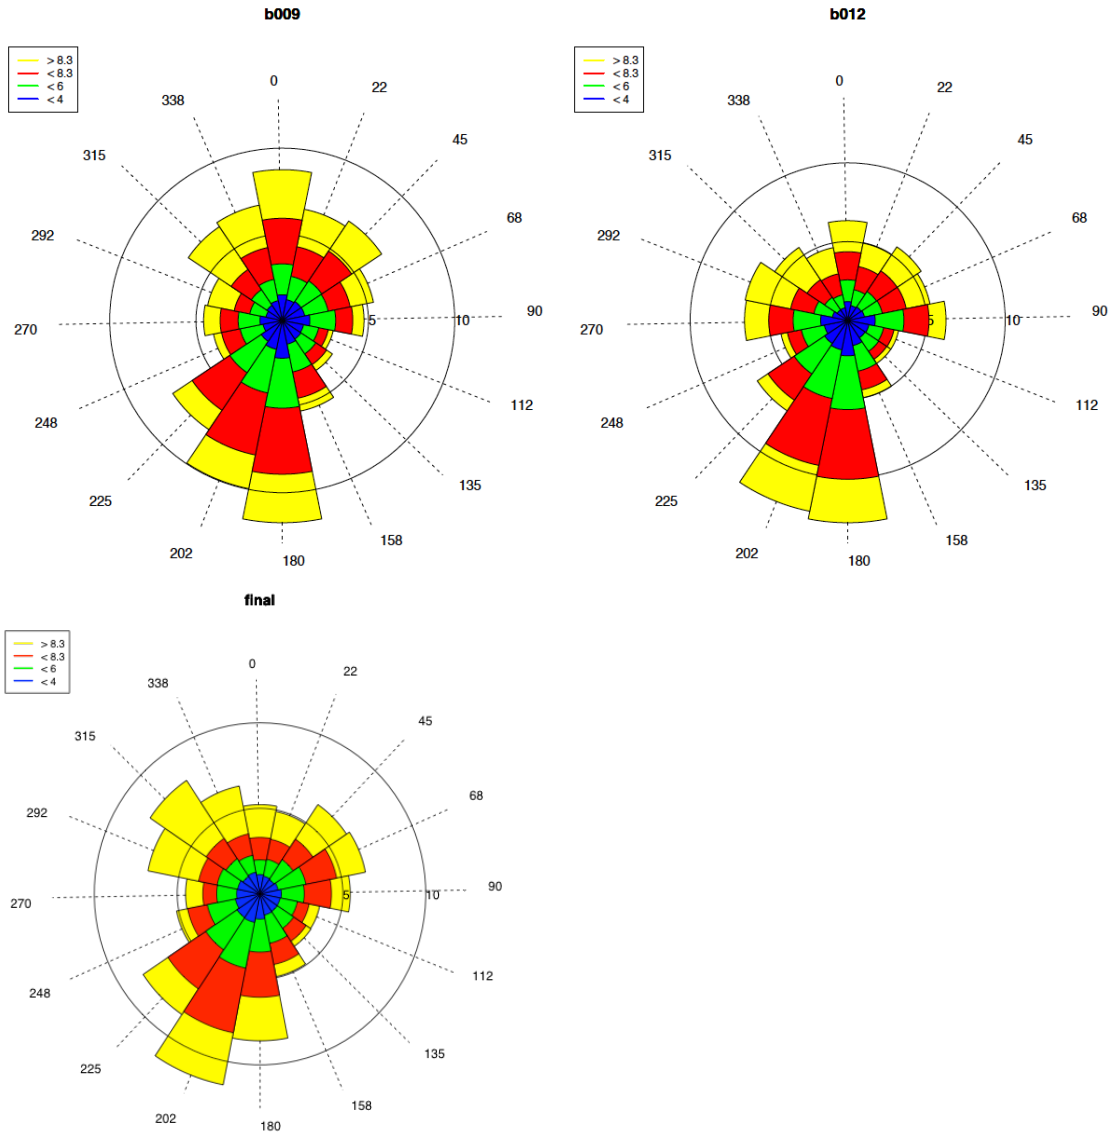


**Figure 18: Wind roses (m/s) for the two buoys near the AOI, and the simulated climatology at the central AOI grid point,.**

One consistent difference between the final and NARR windroses on the one hand, and the buoys on the other, involve the handling of northerly wind directions. The fact that the NARR is showing more winds from the northwest and less from the north compared to buoy observations suggests the possibility that the model is not sufficiently capturing Ekman turning of the low-level winds (Holton, 2004). Ekman turning can occur when a stress occurs across the surface stably stratified fluid, such as the wind blowing above a stable boundary layer. As a result of the stress on the fluid, friction, and the rotation of the Earth, a flow in the stable layer is induced which is slower and to the right of the stress. For a southeastward (northwesterly) wind above a stable boundary layer, this would result is a more southward (northerly) flow in the boundary layer.

### Known artifacts

Transforming data from one geographic coordinate system to another can introduce artifacts of the transformation. In this case we use the term “artifact” to refer to any pattern in a transformed field that is a result of the transformation rather than an underlying pattern that is characteristic of the background field. The temperature fields (air and sea surface) all show artifacts of geographic reprojection (Figure 19). These artifacts are present in the NARR and are not induced by our processing. Instead these unphysical patterns are a product of some geographic reprojection used in creating the NARR. As the temperature artifacts are on the order of 0.1 K and the data will be used mainly in calculating air densities, we are not concerned with the observed artifacts.

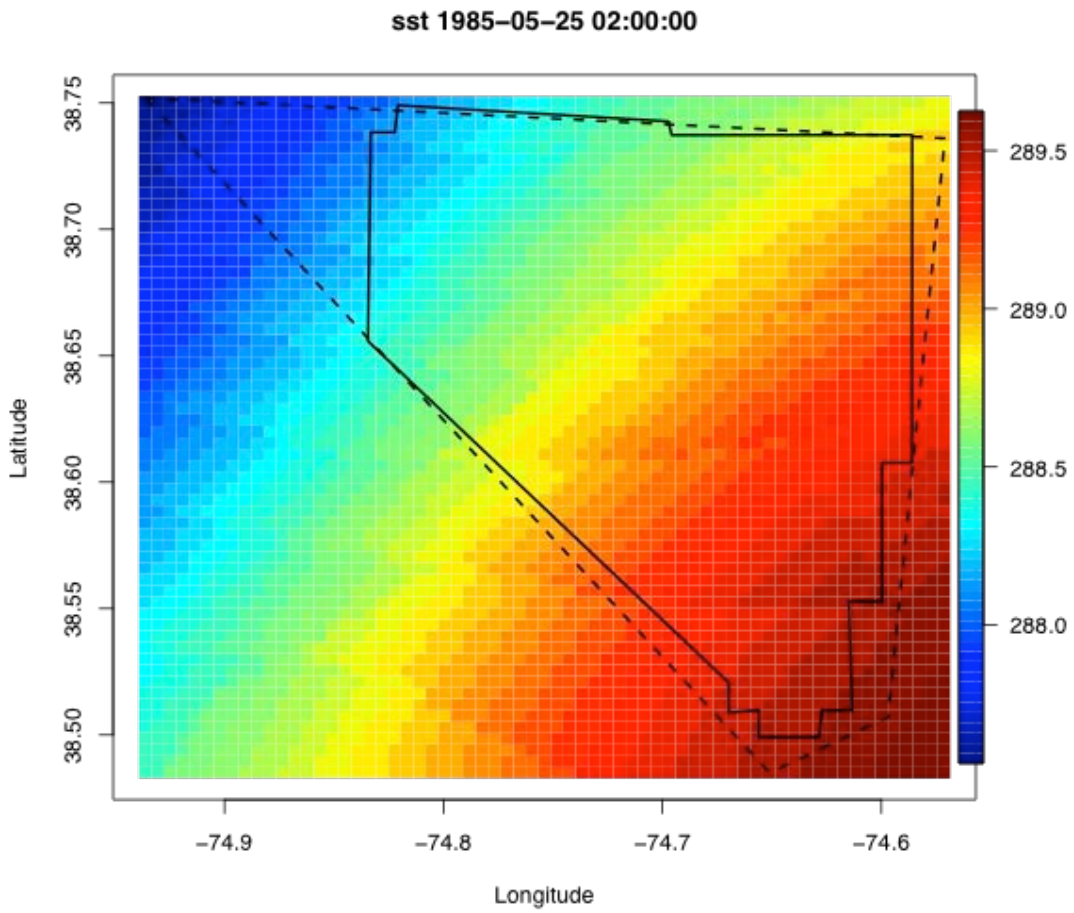


**Figure 19: Map of the sea surface temperatures (K) at an arbitrary time. Note the reprojection artifacts in the NARR dataset.**

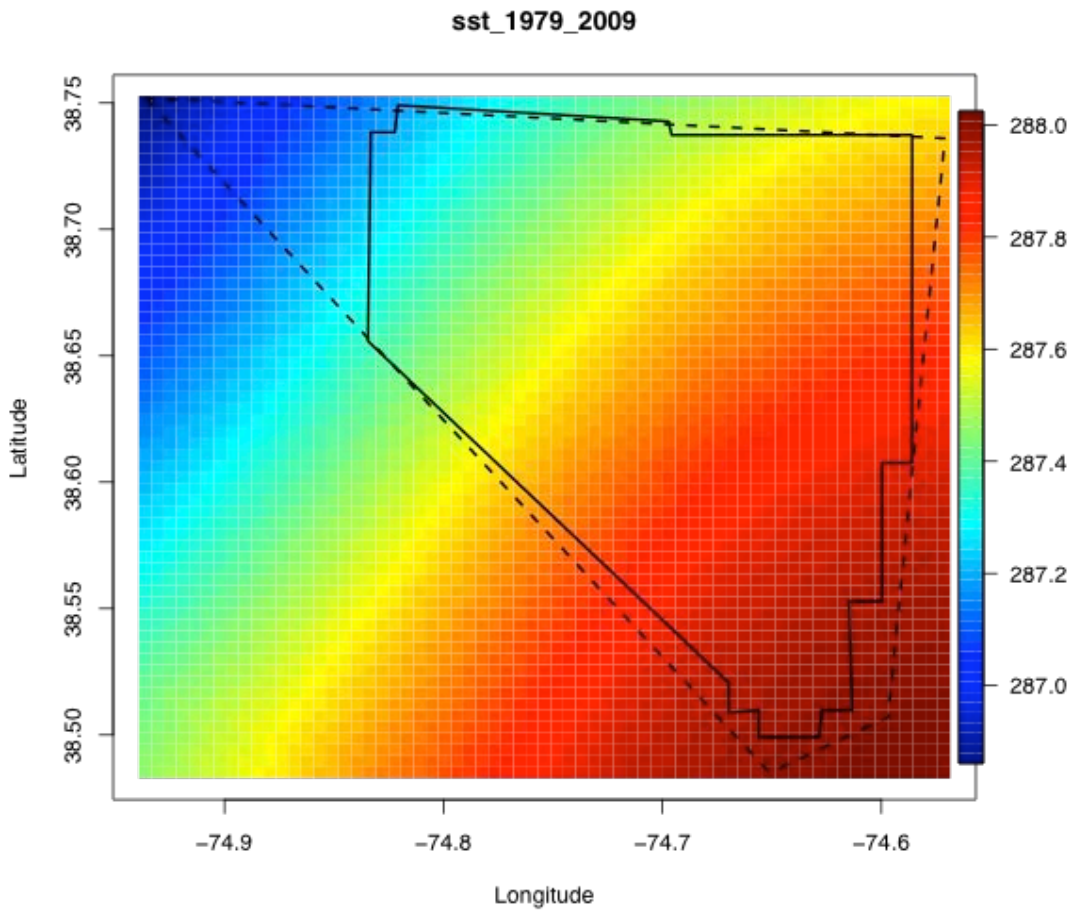


**Figure 20: Map of the annual mean sea surface temperatures (K), averaged over all years in the dataset.**

There are two known artifacts in the NARR wind speeds of which the end user should be aware. A previous wind resource study indicated that the NARR has a bias with lower wind speeds in 2001 and before and higher wind speeds from 2002 onward (Zack et al., 2008). The magnitude of this discontinuity is around 1 m/s and it is limited to the lowest levels of the atmosphere (z ~ 10 m); it is not evident at turbine height. The authors speculate that this is a result of a change in the way that surface wind speed measurements were assimilated into the model. The influence of this change would then drop off with height as the wind field is largely controlled by balloon observations and physical balance. Their study was also conducted in an area of complex terrain in the Pacific Northwest and may not be applicable to offshore regions. If such a problem is contaminating our hub height wind speeds, it is not immediately evident in Figure 10.

There is also a known bug in the code used to generate the NARR data at NCEP which occasional causes bad values for the wind speed at a height of 30 m. This bug is limited to coastal areas which neighboring points having non-negligible terrain. As we are offshore and not near significant terrain, this bug is unlikely to

influence our results. A discussion of this bug is available online.[3] The pressure-level NARR data is calculated in a different method from the above-ground-levels (i.e., 2 m, 10 m, 30 m). Only the NARR pressure levels data, and not the 10 m or 30 m winds, are used by the WRF simulations for its initial conditions so we expect no influences on our high-resolution simulations. WRF chooses to ignore above-ground level for a number of reasons including data quality concerns, the fact that the WRF terrain and NARR terrain do not match due to differing resolution, and a preference for pressure levels which more closely follow the mass-based vertical coordinates in WRF. Winds at fixed heights above the ground in the model output are calculated using a different vertical interpolation scheme within the model than those on the model's vertical levels. This discrepancy can occasionally result in 10 m wind speed values that are anomalously high compared to the higher levels (Figure 21).

[3] http://www.emc.ncep.noaa.gov/mmb/rreanl/faq.html#zero-30m-winds

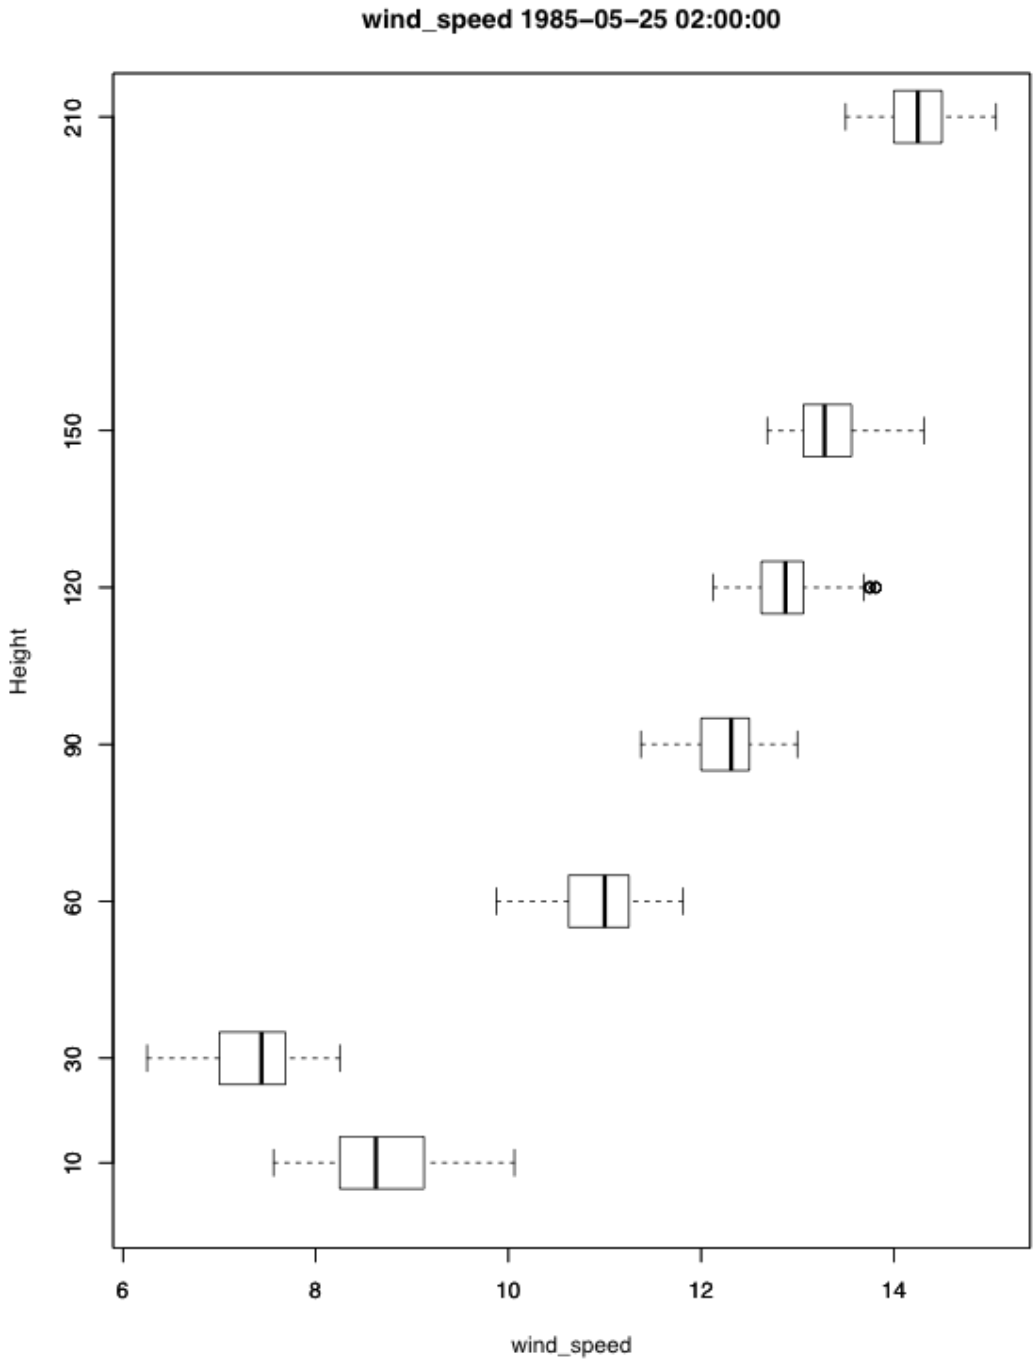


**Figure 21: Box plot of wind speeds (m/s) as a function of height above the surface (m). The box plots show the spatial variability throughout the domain at a single time (labeled in the title). The thick black lines are the median values. See Figure 2 for an explanation of the box plot.**

## D. Conclusions

The time series for Task 2 has been generated on the specified 500 m grid for the AOI at the specified altitudes. The time series spans the period 1979 – 2009, providing a total of 32 years. As there is no fixed-point verification data located inside of the AOI, we instead compare our time series against the wind speed distributions from two nearby buoys. Of these buoys, only one was used in the model sensitivity analysis performed in Task 1. The use of only one buoy in the calibration and two in the validation is far optimal, but was the best available. The limited observations used in this study should be kept in mind when examining the final time series. The buoy distributions suggest that our time series may have a high bias. After applying the bias correction for the observations height that is documented in the Task 1 Final Report, the corrected bias is still positive, but less than the observation error 1 m/s. This high bias could also be the result of the 10 m wind speeds being calculated using a different method than the pressure-level data which was used to initialize the WRF simulation. Artifacts of this difference are evident in Figure 21. We expect that the error in our wind speed distributions is likewise approximately 1 m/s.

The final time series should be usable to estimate electricity generation for a proposed wind farm. When producing the estimates, caution should be used when accounting for so-called “ramp events” when power production changes rapidly. Ramp events are important to understand for the integration of wind power into the power grid. In case of a down ramp (sudden decrease in wind power production), a utility must be ready to increase production from other sources like fossil fuels. This is more of a operational concern for day to day integration of wind power in the grid than it is a concern for a resource study. As long as a ramp event is properly forecast the utility can plan on compensating for it by adjusting generation from other sources. The dependability and flexibility of the power generation may influence the price paid to the power generator and so this may need to be considered in an economic model. Due to the sampling methodology used in this study the time series may not be cleanly continuous across the three hour intervals bounded at 00, 03, 06, …, 21 UTC. The background flow, which is driven by the NARR, will vary smoothly, but the perturbations of the WRF may produce artifacts that influence the prevalence of ramp events. A spectral analysis (in time) should be used to determine the extent to which these artifacts are present, and whether a spectral filter (see, e.g., Turner et al. 2011) should be applied for a particular application. It may be desirable to account for the artificial presence of excess ramp events in an economic valuation.

# Appendix

## Location of vertical levels

In total we use 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", $\eta$, defined as:

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

Where, $p$ is the pressure at level $\eta$, 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 2 and shown in Figure 22.

**Table 2: Vertical $\eta$ 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 |

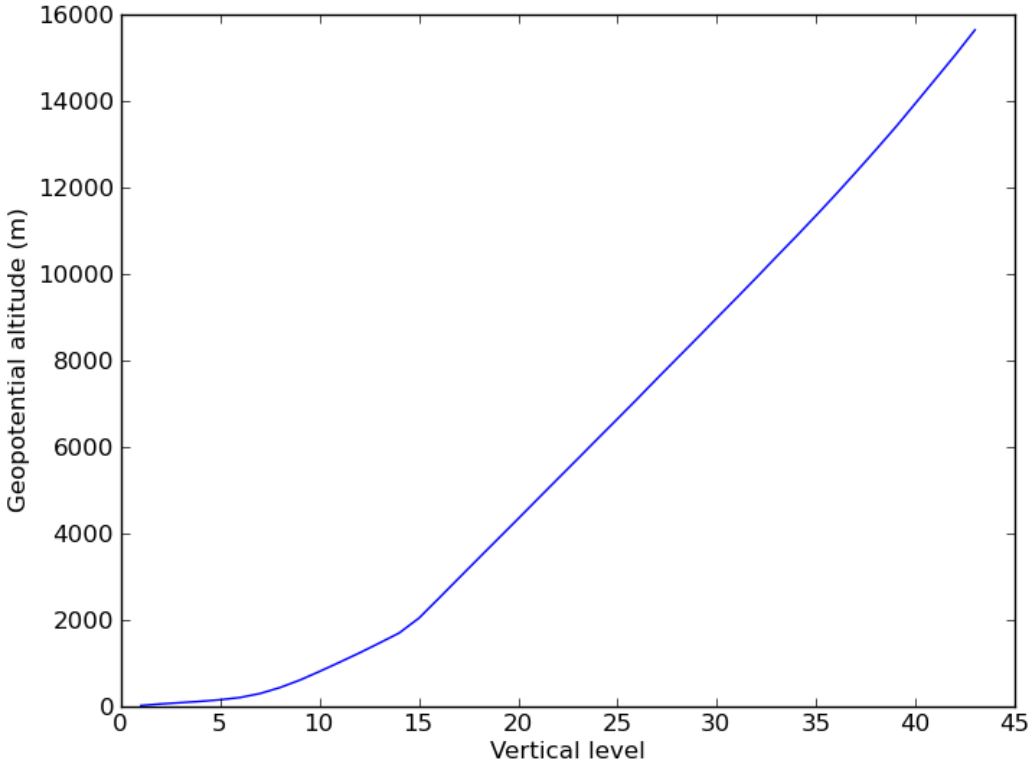


**Figure 22: Distribution of vertical levels.**

## 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 common practice 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 ½ 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).

## NetCDF 4 file format

The final time series is delivered in network Common Data Format(netCDF) 4 files with one file per month. NetCDF4 is maintained by Unidata group, which, like NCAR, is a member of a University Corporation for Atmospheric Research (UCAR). Software and documentation are available online at: http://www.unidata.ucar.edu/software/netcdf/. NetCDF4 is commonly used for weather and climate storage, and is essentially the same as the HDF5 format which is more popular with satellite data.

The data was encoded using the Python bindings to the netCDF version 4.1.3 library on a Macintosh server-grade computer. As netCDF4 is a self-describing format, all relevant metadata including variable names, units, and coordinates are saved within the files. All variables are saved on the fixed latitude-longitude grid with the relevant dimensions ordered: (t-time, z-height, y-latitude, x-longitude). Libraries are available in the C, C++, and Fortran languages. Bindings are available to a number of other languages including but not limited to Python, R, IDL, and MATLAB. Reading the data was tested using both Python and R on Macintosh and Linux machines to confirm inter-platform compatibility.

## Assumptions and limitations

Throughout this we report we have presented a number of assumptions as they were encountered. Here we summarize these limitations for the reference of the reader.

- All of the model calibration and validation has used two offshore buoys that are nearby but outside of the AOI. These observations were taken at a height of 5 m whereas the final deliverable focuses on altitudes in the tens of meters.
- The numerical modeling has been performed at a much higher resolution that the data used to initialize the simulation. While the model physics will allow for process to develop at the finer scale, those process will have to be the result of the model physics and not driven by observations that were included in the initial reanalysis.
- The use of coarser initial conditions means that the model are simulating only the details of the large scale winds. This can result in apparent large errors when judged by traditional metrics like root-mean-square error (RMSE) when a forecast contains phase errors (i.e., errors in timing of a frontal passage even when the magnitude is correct).
- Model levels specified using mass-based coordinates for numerical reasons. This requires the use of vertical interpolation to output the final data on constant altitude surfaces.
- Some small errors are introduced by the reprojection and interpolations of the NARR from a Lambert conformal projection on pressure levels to the model grid.

- A boundary layer parameterization was used to represent the transfer of momentum between the surface and atmosphere as the resolution of the simulations was too coarse to explicitly resolve boundary layer eddies.
- Due to the sampling strategy employed, we expect discontinuities between the 3-hour NARR intervals. This is not expected to influence the total wind energy generation, but will need to be accounted for in any economic models considering the influence of ramp events on integration into the power grid.

# Appendix: Hurricane Wind Assessment Offshore Delaware

This Appendix summarizes the hurricane wind assessment done for the area of interest (AOI) offshore Delaware. The figure below shows the area of interest as a blue polygon overlaid with historic hurricane tracks from 1851 to 2010. Note that this is a very wide area shown and later figures zoom in into the AOI in order to show hurricanes that would be relevant to the AOI. Note that the lease area is not easily visible and lies within the blue oval in the figure below. The brown rectangle in the figure below indicates the area that is zoomed in later figures to show additional detail on hurricane tracks.

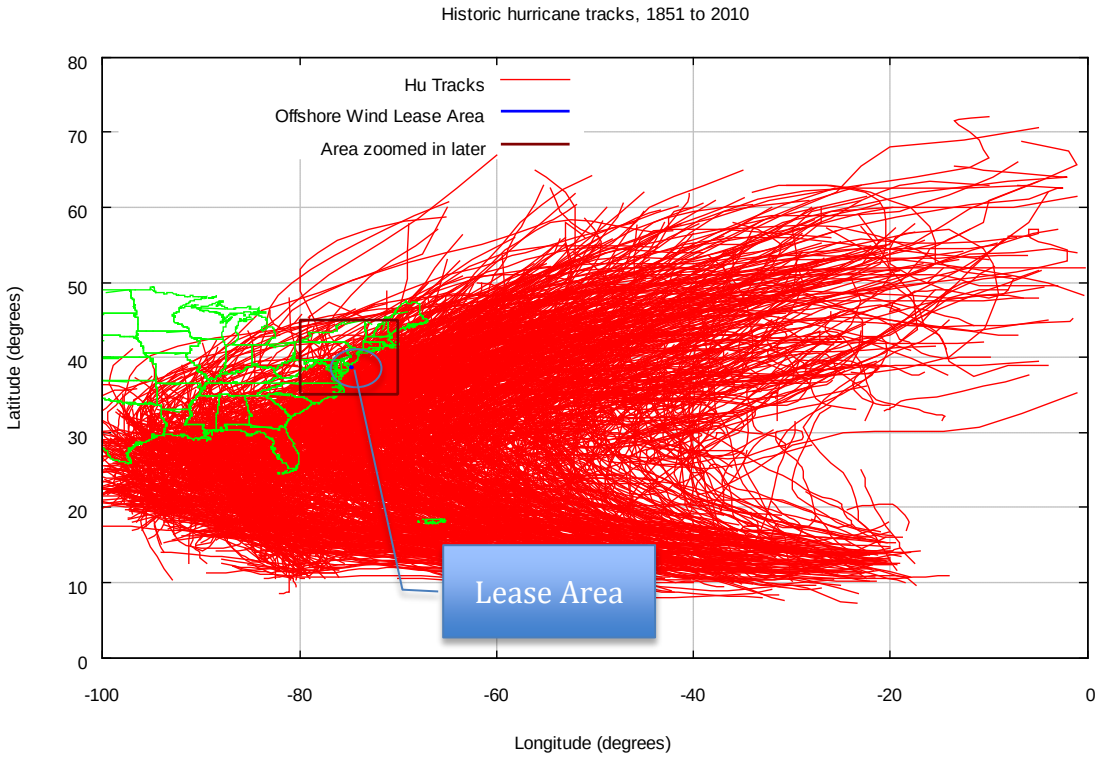


**Figure 23: Area of interest overlaid with historical tracks for past 160 years. Note that only country boundary for USA is shown to convey scale of map shown.**

The next six figures show the area of interest with the maps being zoomed in closer to the lease area (the USA land boundary is not shown). The zoomed in area is overlaid with hurricanes tracks of increasing severity, and as one can anticipate the numbers of hurricanes become sparse as the severity of hurricane increases. Note that these hurricane strength categories are defined by NOAA when the hurricane makes landfall and the hurricane strength may be different that at landfall when the

hurricane passes near the AOI. It is interesting to note that there are two Category 5 hurricanes (defined at landfall) that passed close to the AOI, offshore Delaware.

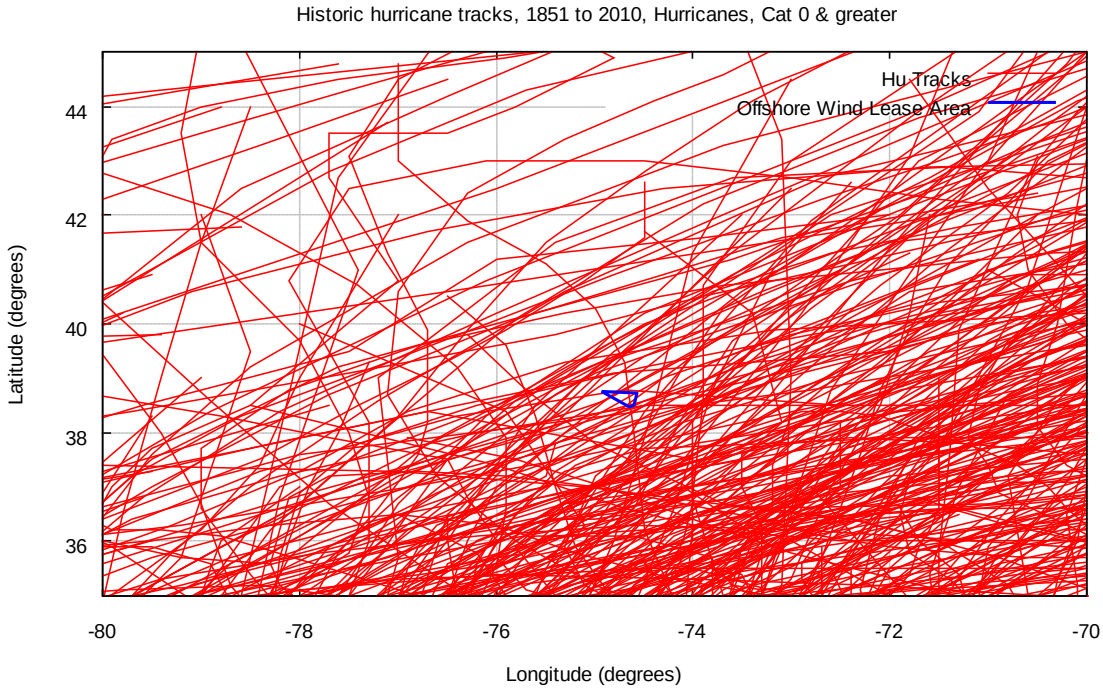


**Figure 24: Tracks for historic storms with landfall category of 0 or greater (i.e., all storms in the NOAA hurricane dataset)**

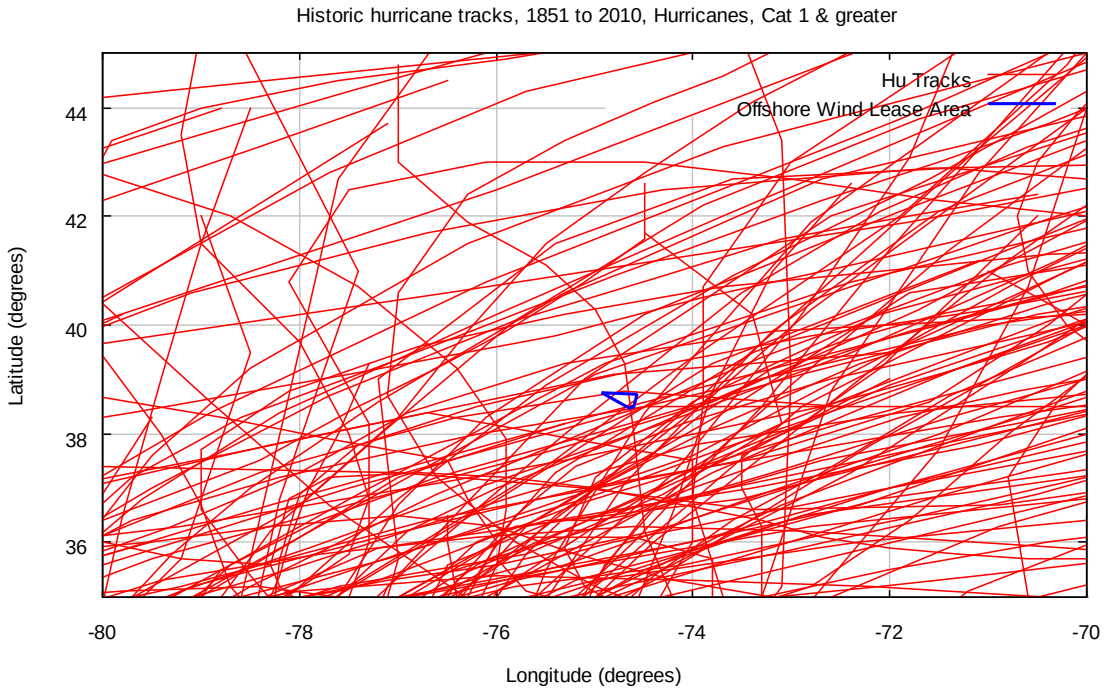


**Figure 25: Tracks for historic storms with landfall category of 1 or greater**

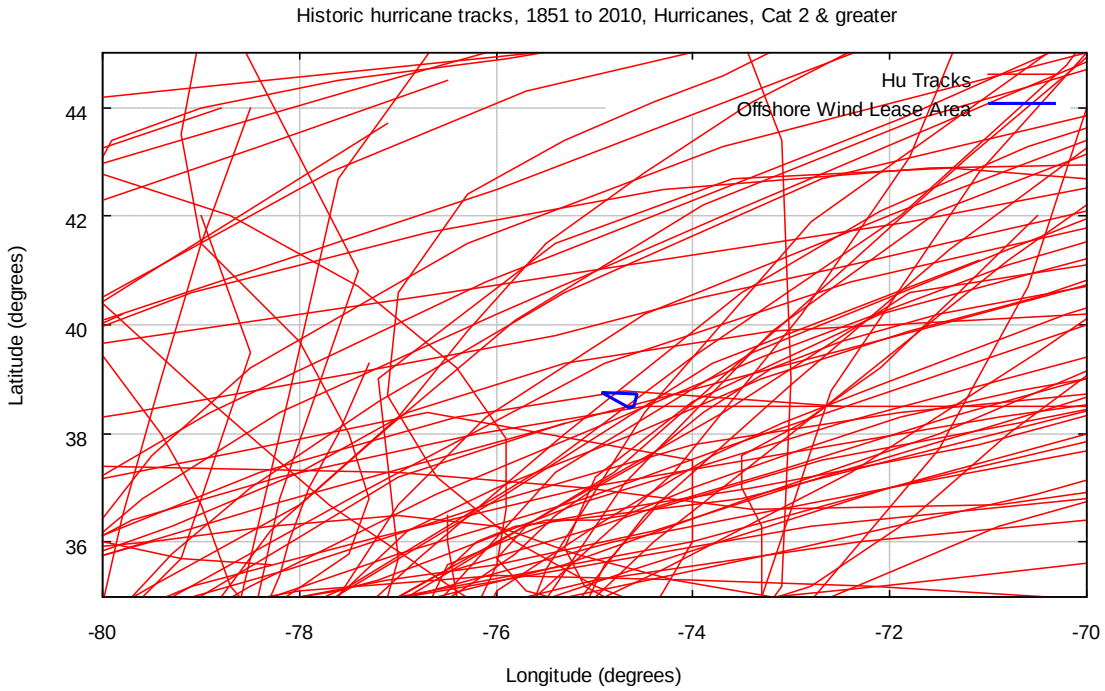


**Figure 26: Tracks for historic storms with landfall category of 2 or greater**

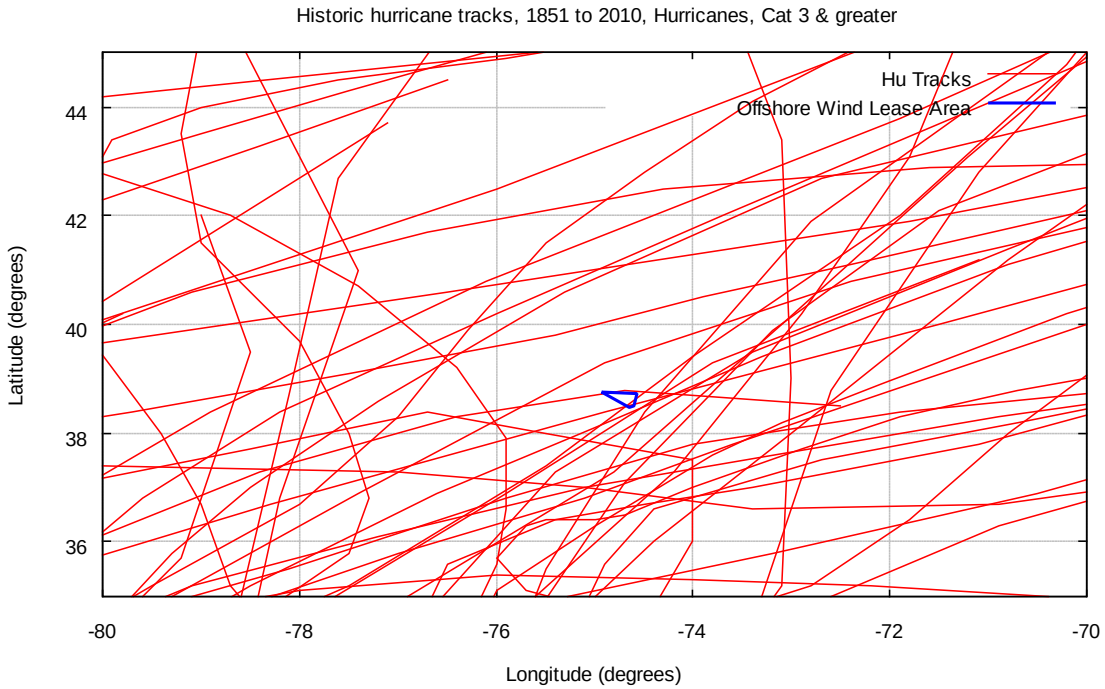


**Figure 27: Tracks for historic storms with landfall category of 3 or greater**

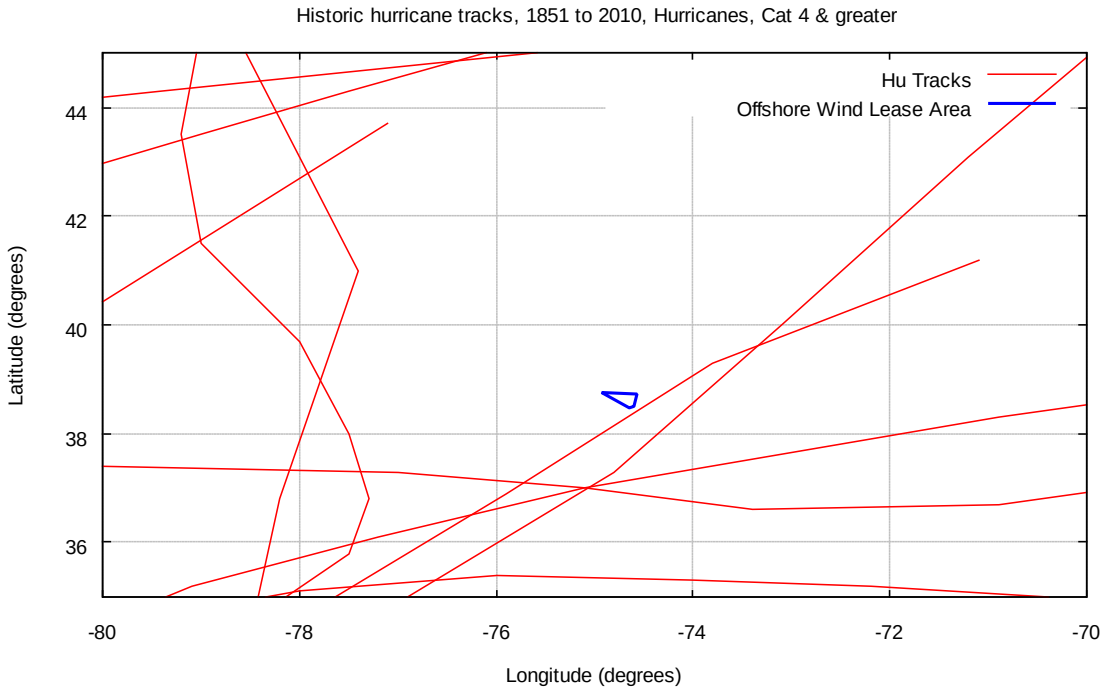


Figure 28: Tracks for historic storms with landfall category of 4 or greater

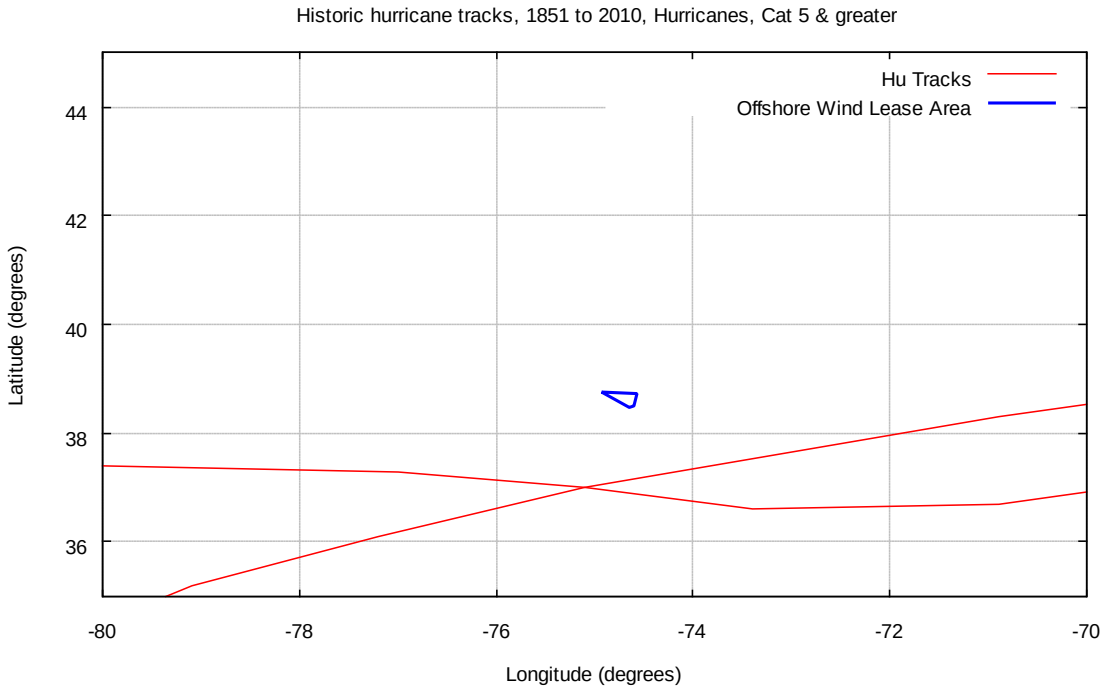


Figure 29: Tracks for historic storms with landfall category of 5 or greater

The objective of hurricane wind speed assessment is more relevant to offshore turbine design strength assessment, than for assessment of electricity generated from non-hurricane winds. Hurricane winds are rare in nature; however, the relative frequency varies along the coast with greater hurricane frequency in lower latitudes along the east coast of the US. There are various methods that can be pursued for robust assessment of hurricane wind speeds. A method is presented here based only on historical hurricane analysis. Other methods utilize simulation of synthetic hurricanes to fill in gaps that historical hurricanes may not have captured thus far. We use only the historic hurricane in this study to demonstrate the methodology and to convey the order of magnitude of hurricane winds that are possible in the lease area. The specific steps adopted in this study are discussed next, with the end objective to develop estimates of wind speeds are various extreme return periods (e.g., 100-years), meaning, wind speeds are expected to exceed once every 100 years, or in other words, wind speeds that are expected to occur with a probability of 0.01 or lower in any given year.

We gridded the AOI with a roughly 500meter grid to develop hurricane wind estimates for each historic hurricane in the NOAA database (hurdat). This database was downloaded from the NOAA website as of January 2012, and is assumed to reflect the latest hurricanes in the database (which includes hurricanes in 2010). The following steps were used to develop hurricane wind speeds at each grid point.

- Each hurricane in the historic dataset was passed by the grid
- The track time resolution by NOAA is 6 hours, we linearly interpolated this track to generate track points at every 15minutes (to get a better sampling of the hurricane points in time)
- For each grid point with the figure shown below, a distance to the hurricane track point and the wind speed from that track point is calculated using formulas mentioned below (taken from metocean literature on this topic). There are more sophisticated models that can employed, for example, those include the “Holland B Parameter” which models the left to right asymmetry of hurricanes, however, this sophistication was judged as unnecessary for the objectives for this study, since we intend to convey only the order of magnitude of extreme wind speeds in this study. A sample wind time history at a grid point is shown in Figure 28. This figure show how the wind speed changes with time at a grid point of interest as the hurricane passes by the grid point. The wind speed increases gradually and then reduces as the hurricane moves away from the grid point. Note only an average wind speed is captured in this study, in reality, there is uncertainty about this mean wind speed owing to turbulence, which again is not captured in this study since it is not the focus here.
- As one hurricane passes by, the maximum wind speed generated by this hurricane at each grid point is saved (e.g., the largest wind speed shown in

Figure 28). This maximum wind speed is considered as the hurricane wind for the hurricane just analyzed.

- Similarly, all hurricanes in the historic database are analyzed for every grid point, to generate a vector of wind speeds at each grid point, with the size of the vector being the number of hurricanes analyzed. Note that distant hurricanes have little impact on grid points of interest.
- For each grid, the vector of wind speeds is used to develop an annual probability distribution of wind speeds. This distribution is then used to read the wind speeds for specific return periods (see Figure 29). For example, the 100-year winds are equal to wind speeds that have an annual exceedance probability of about 0.01 (=1/100).

Validation of the hurricane wind speeds was done via comparisons to ASCE 7 wind design guidelines basis[4]. This paper has wind maps for 100-year that shows a wind speed (3-second gust at 10meter height) of about 45 to 49 meters per second near the area of interest. This when adjusted to a 1-minute wind speed translates to a wind speed range of 39 to 42 meters per second. This compares closely with 100-year wind speed in the area of interest, which is about 43 meters per second (see Figure 29). This agreement is remarkable given the potential differences in methodologies used between this study and the ASCE study.

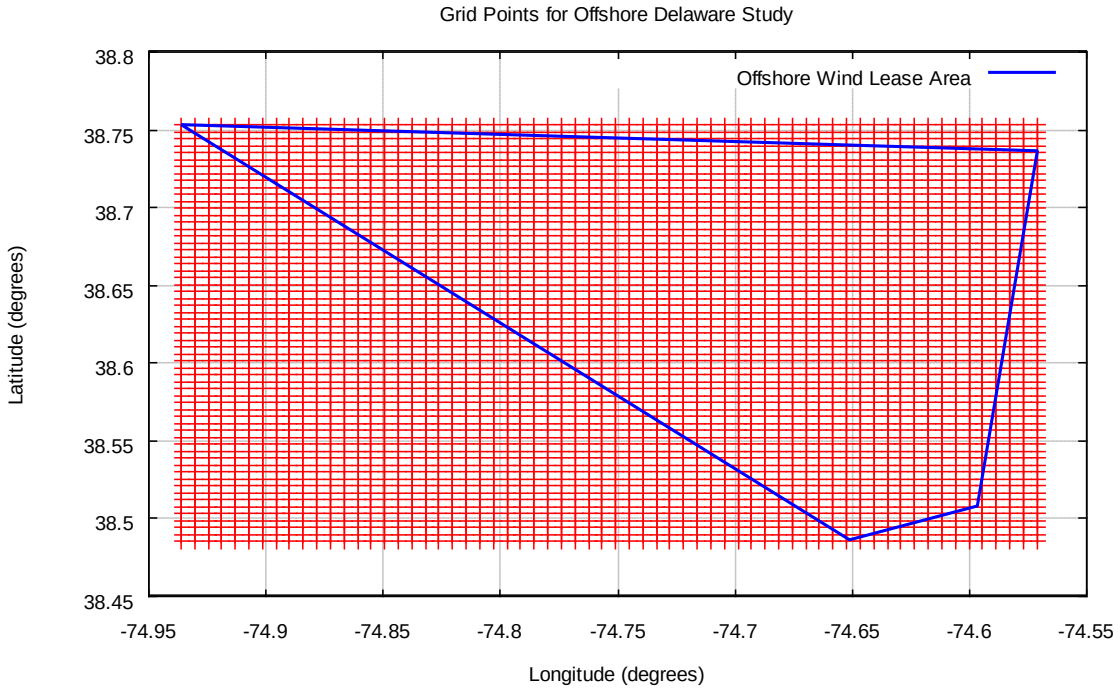


**Figure 30: Area of interest, offshore Delaware, and grid points developed to analyzed hurricane winds**

[4] Vickery et al. Ultimate Wind Load Design Gust Wind Speeds in the United States for Use in ASCE-7. Journal of Structural Engineering. ASCE. MAY 2010

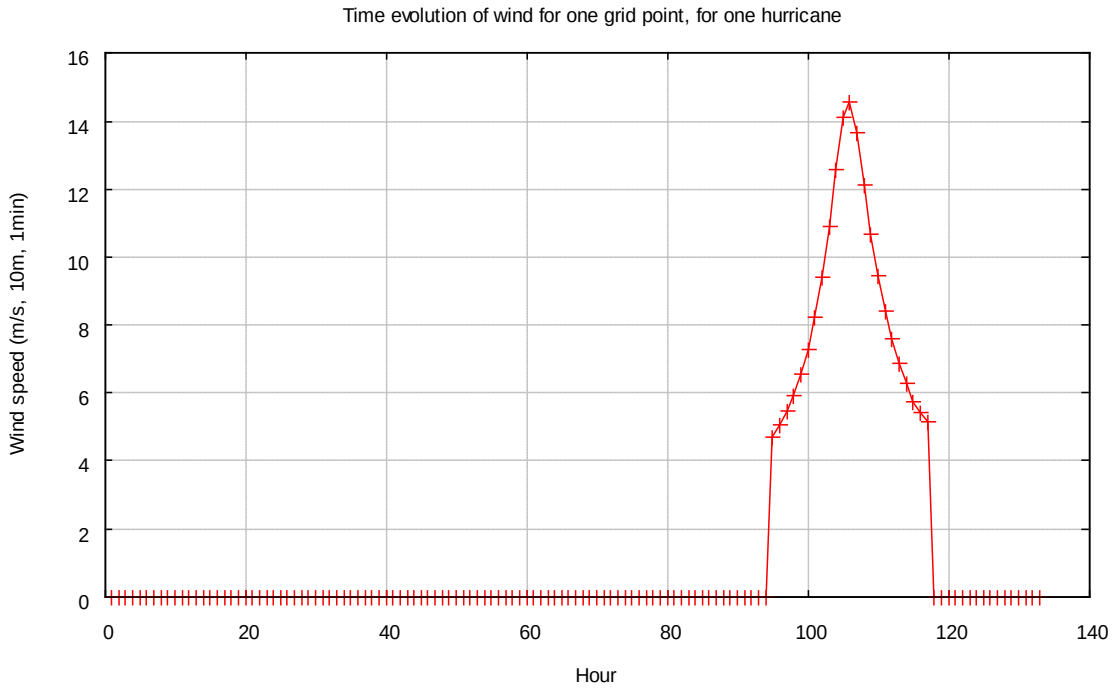


**Figure 31: Time history of hurricane wind speed (mean) as hurricane passes by a given grid point within area of interest**

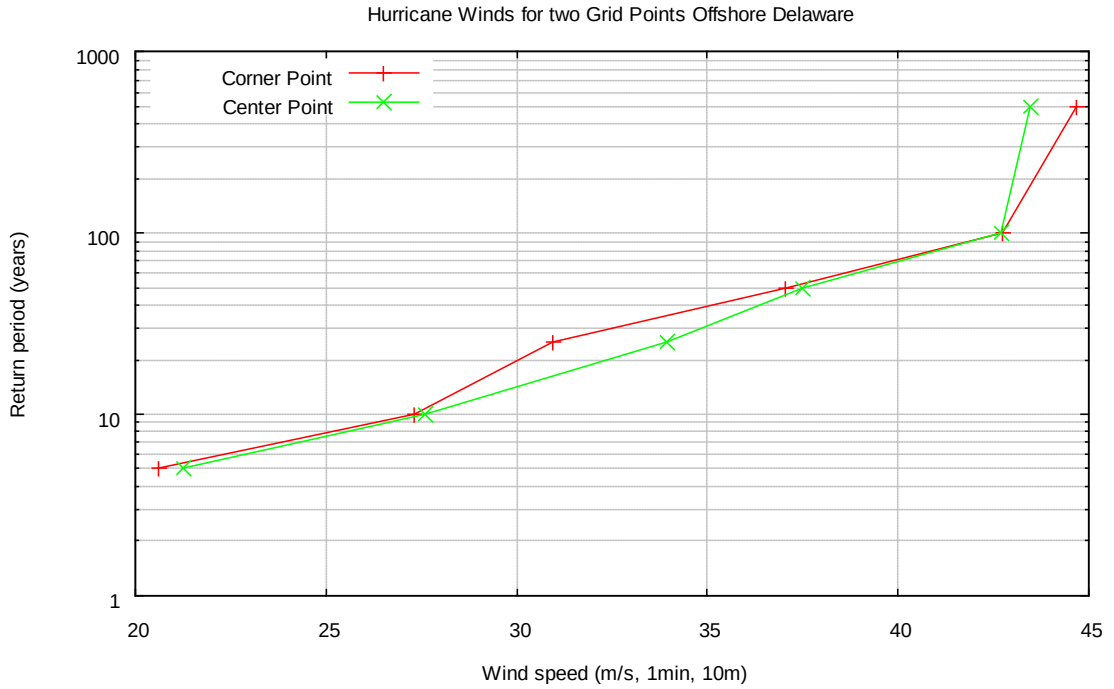


**Figure 32: Hurricane wind speed for different return periods for two sample grid points within area of interest**

The next three figures show the 100-, 10-, and 500-year return period wind speeds caused by hurricanes for the area of interest. As orientation on the figures, the surface shown using red gridded line are the wind speed values (the vertical axis) at each grid point within the latitude and longitude range shown. There are contour lines shown on the surface as well as on the x-y plane; these contour lines are specific wind speed values over the area of interest. For example, the wind speed value of 45 m/s, shown as a green line, occurs in one corner of the gridded region. These contours are being shown to easily see the wind speed values in a 3-D chart.

Given the simplicity of the model and given the uncertainty in the parameters involved, the wind speed for a given return period is assumed here to be similar (constant) over all the grid points in the area of interest.

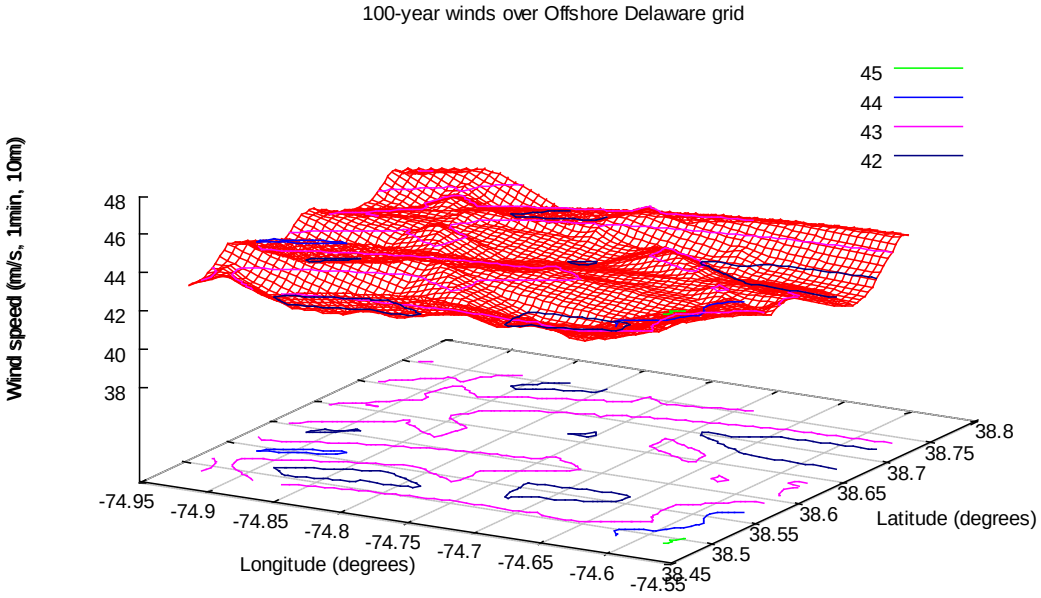


**Figure 33: 100-year wind speeds from hurricanes over grid points in area of interest**

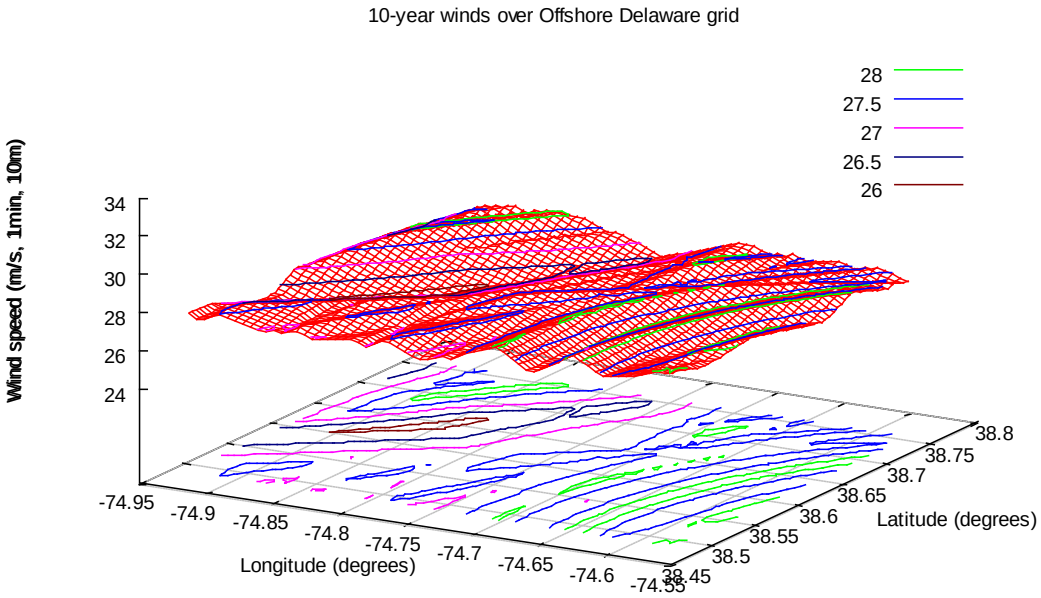


**Figure 34: 10-year wind speeds from hurricanes over grid points in area of interest**

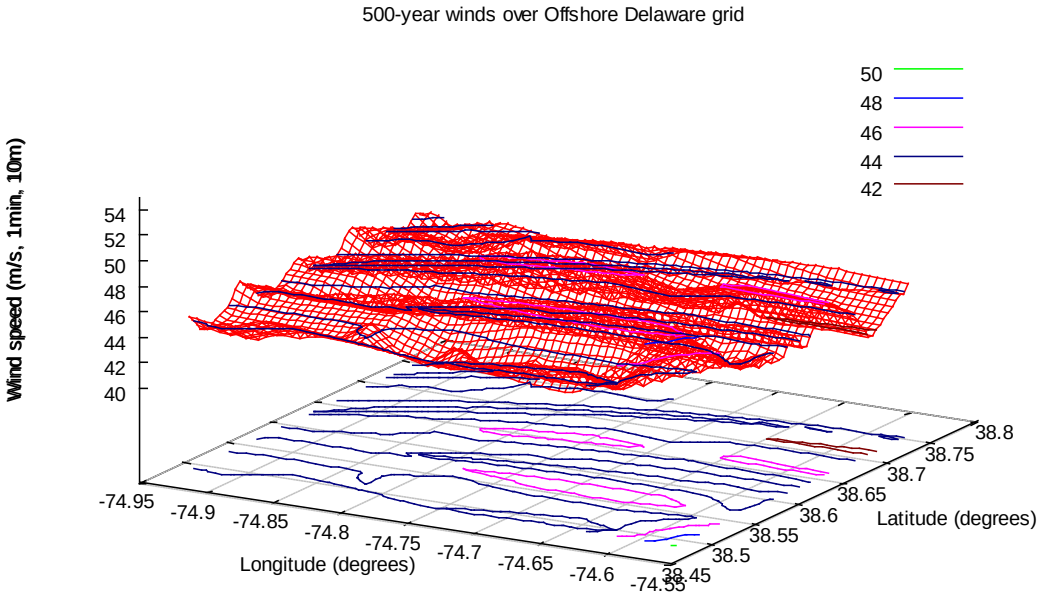


**Figure 35: 500-year wind speeds from hurricanes over grid points in area of interest**

## List of Acronyms and Abbreviations

- AFWA – U.S. Air Force Weather Agency
- AOI – area of interest
- ARW – Advanced Research WRF
- ASOS – Automated Surface Observing System, the official NWS weather station network
- BOEM – Bureau of Ocean Energy Management
- MYJ – Mellor-Yamada-Janjic 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

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

Evenden, G.I., 1990: Cartographic projection procedures for the UNIX environment – A User's Manual. *U.S. Geological Survey Open-File Report 90-284*, 68 p.

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.

Holton, J.R., 2004: "Chapter 5 – The Planetary Boundary Layer". *Dynamic Meteorology*. International Geophysics Series. **88** (4th ed.). Burlington, MA: Elsevier Academic Press. pp. 129–130.

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., 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.

MMI Engineering and Atmospheric and Environmental Research, Inc., 2011: Prediction of Wind Energy Resources on the Outer Continental Shelf with Weather Models. BOEM Tech. Rep. M11PC00017, 57pp.

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.

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

Schwartz, M., D. Heimiller, S. Haymes, and W. Musial, 2010: Assessment of offshore wind energy resources for the United States. NREL/TP-500-45889, National Renewable Energy Laboratory. http://www.nrel.gov/docs/fy10osti/45889.pdf.

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.

Smirnova, T. G., J. M. Brown, S. G. Benjamin, and D. Kim , 2000: Parameterization of coldseason processes in the MAPS land-surface scheme. *J. Geophys. Res.*, **105**, 4077-4086.

Turner, R. et al., 2011: Creating Synthetic Wind Speed Time Series for 15 New Zealand Wind Farms, *J. Appl. Met Cli.*, **50**, 2394-2409.

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.

Zack, J. W., K. T. Waight, S. Young, G. E. Van Knowe, and M. Bower, 1998: Estimation of monthly and annual local wind speed anomalies through the use of high-resolution numerical simulations. *17th Conf. on Applied Climatology*, Whistler, BC, Amer. Meteor. Soc., 10.3.