Fast Verification of Wind Turbine Power Curves: Summary of Project Results

by:

Cameron Brown – s101246

Supervisors:

Pierre Pinson, DTU, IMM

Mike Courtney, DTU, Wind Energy

Technical University of Denmark

Kongens Lyngby

July 13, 2012

IMM-MSC-2012-72

Technical University of Denmark
Informatics and Mathematical Modeling
Building 321. DK-2800 Kongens Lyngby, Denmark
Phone +45 45253351, Fax +45 45882673
reception@imm.dtu.dk
www.imm.dtu.dk

Executive Summary

A new method has been developed for analyzing power performance measurements using a stochastic differential equation called the Langevin equation on high frequency wind turbine measurement data sampled at one sample per second or more. The aim of this method is to provide an alternative to the current industry standard described in IEC 61400-12-1. This new method was implemented and analyzed in a number of test cases to evaluate its effectiveness and robustness as a part of the FastWind project at the Technical University of Denmark's (DTU) Wind Energy department. The FastWind project aims to develop and test this method for use in the Danish wind energy industry.

This new method was successfully implemented for use in this project as well as future work in the FastWind project. Several methods of estimating the power curve uncertainty were developed and implemented; including an analytical equation and two data resampling methods.

Using measurement data from DTU's Nordtank wind turbine at the Risø site, the practical application of this new method was tested. It was found that the method is sensitive to some of the input parameters chosen by the user. Additionally, it was found to be sensitive to the coherence of high frequency wind speed and power measurements. These sensitivities indicate that the base method is not ready for application in industry but the method needs some further development to improve the robustness of the method's results. However, potential solutions to these difficulties were proposed and show some promise for improving the outputs of the method.

Proponents of this new method claim that the method for effectively uses measurement data and therefore will require less data, and therefore shorter measurement campaigns, to obtain a power curve with a desired level of certainty. Additionally, it has been claimed that the method is insensitive to turbulence intensity and site specific conditions; unlike the standard IEC method. The potential benefits of this new method were tested and some positive results were obtained. The method was shown that it may converge faster than the standard IEC method requiring less data to obtain the same level of certainty. Furthermore, the method shows some insensitivity to turbulence intensity indicating that site specific turbulence will not affect the power curve obtained. Although this new method has some challenges associated with it, they are not unsolvable and these benefits show that the method may be worth the time invested to overcome these difficulties.

Table of Contents

1 Introduction

This report summarizes the work done by Cameron Brown during his thesis project as a part of the Master of Science in Wind Energy degree at the Technical University of Denmark (DTU). This project has been completed in collaboration with the Wind Energy, and the Informatics and Mathematical Modeling departments at DTU. The thesis project was a component of "FastWind"; a larger research project at DTU Wind Energy with collaboration from Vestas Wind Systems.

The goal of the FastWind project is to more effectively use wind turbine measurements which are typically recorded at several samples per second. Due to fluctuations in the measurements caused by phenomena such as atmospheric turbulence, the data can be difficult to interpret so it is typically reduced to 10 minute mean values during analysis. The project aims to utilize the potentially useful information contained in the raw high frequency data that is lost during averaging. This potentially allows for faster optimization of turbine design and control strategies at no added cost by simply developing and implementing new methods to analyze this high frequency data.

One of the classic examples of analyses using 10 minute mean data is the analysis from the International Electrotechnical Commission (IEC) wind turbine standard 61400-12-1 regarding power performance measurements of electricity producing wind turbines [1]. A new method which attempts to utilize the high frequency data for measuring the power performance of a wind turbine has been developed at the University of Oldenburg [2]. This method uses the stochastic differential equation called the Langevin equation to model the power output of a wind turbine under the fluctuating driving force of turbulent wind. One of the goals of the FastWind project is to develop a familiarity with this method, verify its effectiveness, and promote its proliferation within the Danish wind energy industry. For simplicity, these two methods will be referred to as the "IEC method" and the "Langevin method".

The objective of this thesis project was not to improve the method developed in Oldenburg but to test the practical application of the method as is. The individual goals of this thesis project were summarized as follows:

- Since this project was performed at the beginning of the FastWind project, the first goal was to gain an understanding of the method developed in Oldenburg. The method was then to be implemented for use in this thesis and further work during the remainder of the FastWind project. The method was then to be documented such that it could be understood, and be reproduced or improved during further work in the FastWind project. The base method and its implementation are described in Section 4.
- Secondly, practical implementation of this Langevin method was investigated to see if the method requires more development before it can be applied in industry. Tests were performed to determine how robust the method's outputs are to inputs parameters and data pre-processing chosen by the user. For example, pre-averaging of the wind speed has been suggested in previous work [3]. Methods for choosing these inputs and the necessary pre-processing were investigated and are discussed in Section 5.
- Finally, the effectiveness of the Langevin method was examined. Some of the claimed benefits over the traditional IEC method made by the inventors of this Langevin method were investigated. The Langevin method was compared to the IEC method in a number of tests as described in Section 6.

Before discussing the implementation of the Langevin method and the results, a background on the difficulties with power curve measurements and some of the current solutions addressing these difficulties are discussed in Section 2. The measurement data used for the testing the Langevin method is then described in Section 3.

2 Background on Power Curve Measurement

Before describing the Langevin method in detail and the results obtained during this thesis, some background information regarding power curve measurements will be discussed. A Wind turbine's power curve is one of its most important selling features and the reasons for this are described in Section 2.1. However, a power curve is not trivial to measure as is described in Section 2.2. The most common method currently used for determining a power curve, IEC 61400-12-1 is described in Section 2.3 while the proposed Langevin method under investigation in this thesis is introduced in Section 2.4.

2.1 Purpose of a Power Curve

A wind turbine power curve is the relationship between the driving wind speed and the electrical power produced by the wind turbine. The power curve is used by wind farm developers along with the wind conditions at a given site to determine if construction of a wind farm will be profitable. The power curve is also used later after construction of the wind farm for forecasting the wind farms power production helping the wind farm operator to sell the electricity. If the future power production of a wind turbine can be accurately predicted, the risk involved with building and operating a wind farm is decreased; this results in an attractive investment and the construction of more wind farms. Because companies rely on the power curve to determine the economics of a wind farm, wind turbine manufacturers are forced to guarantee their wind turbine's power curves. Since a turbine manufacturer must provide this legal guarantee on their products, they need to ensure that their turbines will live up to this guaranteed power curve.

During a wind turbine's design phase, the power curve can be predicted using analytical techniques such as Blade Element Momentum (BEM) theory [4]. However, before selling a newly designed wind turbine, the power curve must be empirically validated. Also, as wind turbine manufacturers are testing new components or optimizing their turbines, they need a way to empirically validate their models and determine if the changes they have made are providing positive results. Additionally, during operation of a wind farm, a wind turbine's power curve may evolve over time due to wear of turbine components, dirt buildup on the wind turbine blades, and many other effects. Therefore, an accurate method for experimentally measuring a wind turbine's power curve is required. However, some difficulties exist with power curve measurements as will be discussed next in Section 2.2.

2.2 Difficulties with Power Curve Measurements

A wind turbine's power curve must be validated on a full scale prototype turbine. Since modern wind turbines have rotor diameters of over 100 m, it is not possible validate the turbine's power curve in a controlled environment such as a wind tunnel. Therefore, the main difficulties with wind turbine power curve measurements result from the fact that the driving wind speed cannot be controlled. Prototype wind turbines must be validated at a test site exposed to wind driven by the local weather patterns. Therefore, the turbine is exposed to several natural phenomena which make power curve measurements difficult. Any method of analyzing power performance data must overcome these difficulties when trying to determine a wind turbine's power curve. These difficulties have been summarized in the following sections.

2.2.1 Fluctuations in the Driving Wind

Wind low in the atmosphere is affected by the earth's surface causing the development of a turbulent boundary layer [5]. Wind turbines are producing power from this turbulent wind. This turbulence means that the driving wind speed is constantly fluctuating, causing the wind turbine's power production also to constantly fluctuate as illustrated in Figure 1. Additionally, the power production of a wind turbine cannot respond instantaneously to changes in wind speed. There is some inertia in the wake that creates some time delay between a change in wind speed and

reaching the new steady state power production. Furthermore, since the wind speed is constantly changing, it is unlikely that the turbine will reach a steady state before the wind speed has changed again. Therefore, the wind turbine's power production is constantly chasing the steady state value which is constantly fluctuating with the wind speed.

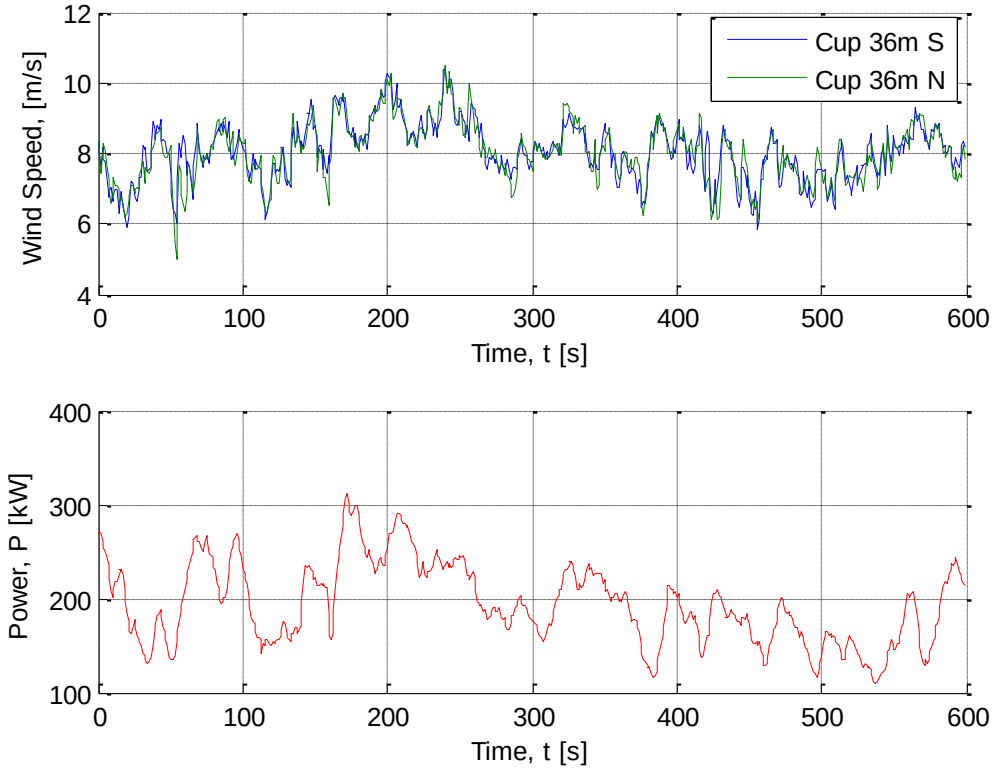

Figure 1 – Example of typical fluctuations in wind speed and power

The steady state power production over the full range of wind speeds are referred to as the "fixed point" power curve or "zero turbulence" power curve. This "fixed point" power curve is what we would ideally extract from the wind speed and power data since this power curve is the most useful for comparison with BEM models and for forecasting of future power production. However, due to the fluctuating driving wind and the inertia of the wake, wind speed and power measurements form a cloud of points around the "fixed point" power curve making it difficult to determine.

2.2.2 Averaging of Wind Speed over Wind Turbine Rotor

The wind speed is not uniform across the wind turbine rotor plane, however the power curve is often described as a function of a single wind speed. The turbulence contained in the boundary layer contains eddies with a range of sizes causing the wind speed to vary over the swept area. Additionally, in the atmospheric boundary layer, the wind speed increases with height, referred to as wind shear. Several functions can be used to describe a typical wind shear profile; one example is the power law profile which is related to the hub height wind speed, $U_H = U(z_H)$, by a constant, α, as shown in Equation (1) [5].

$$U(z) = U_H \cdot \left(\frac{z}{z_H}\right)^{\alpha} \tag{1}$$

Non-uniform inflow causes the kinetic energy flux to vary across the rotor plane. The wind turbine extracts energy from all parts of the swept area and is effectively averaging this kinetic energy flux through the rotor plane. Since the turbine is averaging the energy flux, the power curved could be correlated to a wind speed which is averaged in a similar manner. The wind speed corresponding to

the average energy flux through the rotor plane could be thought of as the "driving" wind speed [6] (referred to as "equivalent" wind speed in [6]; however, this report refers to the same concept as the "driving" wind speed). Using the driving wind speed allows the power curve to be described as a function of a single wind speed which is much more convenient.

An example of a possible wind profile with the equivalent driving wind speed averaged over an example rotor area is illustrated as shown in Figure 2. The hub height wind speed is 10 m/s with a shear exponent, α, of 0.2 and the example rotor has a hub height of 90 m and a diameter of 100 m.

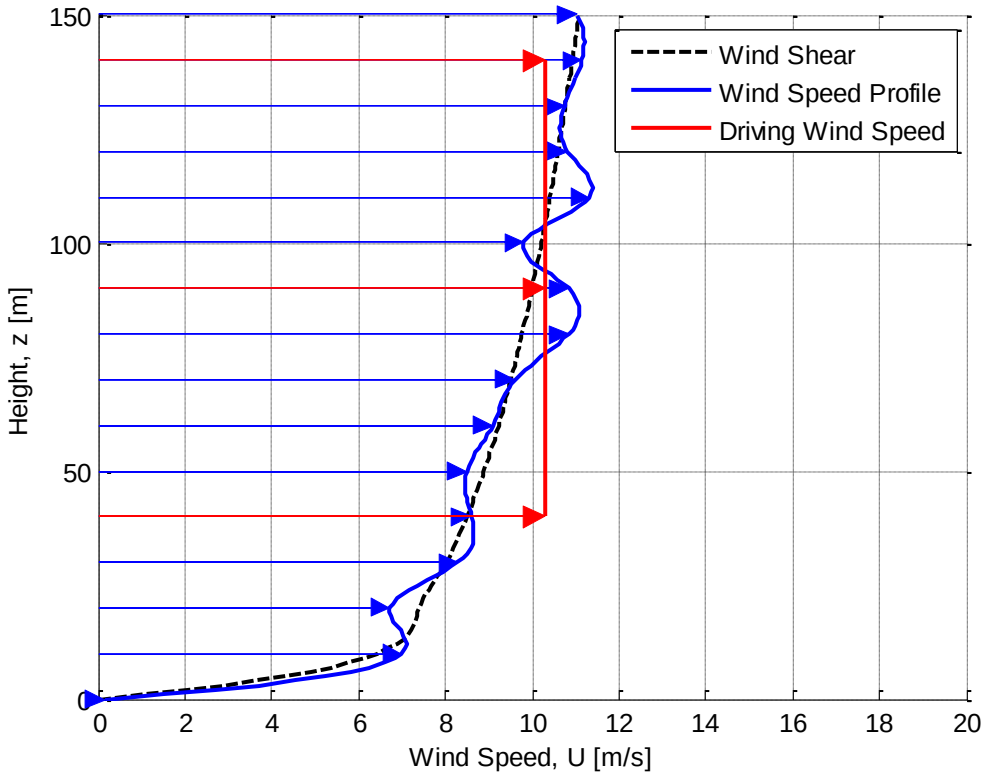

Figure 2 – A typical wind speed profile with the "driving" wind speed over rotor area

Unfortunately, the driving wind speed is difficult to measure since measurement of the wind speed at many points around the swept area would be required followed by some kind of spatial averaging. Obtaining measurements at several points across the rotor plane requires a well instrumented mast which has many cup anemometers at several different heights or a LIDAR; however, these luxuries are not available at all wind farms and test sites.

2.2.3 Spatial Separation of Wind Speed Measurement

The power production measurements should be correlated with the undisturbed wind speed which would be measured at the turbine's rotor if the wind turbine was not present. However, the wind turbine puts a thrust force on the wind in the upstream direction causing it to slow down as it approaches the rotor [4]. Therefore, the wind speed must be measured at a far enough distance that it is not affected by the rotor. Distances of 2 to 4 rotor diameters are accepted according to the IEC 61400-12-1 standard with a recommended separation of 2.5 rotor diameters [1]. For a modern wind turbine with rotor diameters of 100 m and up, this means the met mast must be placed on the order of 250 m or more upstream of the turbine.

In the case that the met mast is directly upstream of the wind turbine as shown in Figure 3 (a), this spatial separation creates a time delay between when the wind speed is measured and when the

wind passes over the rotor. Additionally, the turbulence in the wind field is continuously evolving, so the wind speed measured by the met mast will have changed by the time it reaches the turbine.

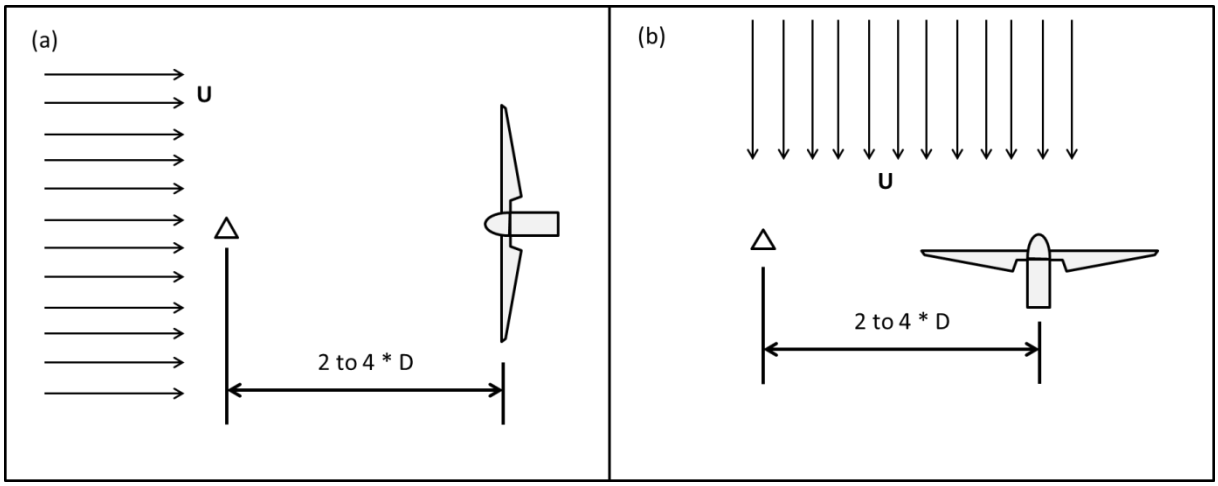

Figure 3 – Illustration of the separation between a wind turbine and a met mast with different incoming wind directions.

In the case when the line between the met mast and turbine are perpendicular to the incoming wind as shown in Figure 3 (b), the wind measured by the met mast will not pass over the turbine. Differences in the terrain upstream of the met mast compared with the terrain upstream of the wind turbine could cause differences in the behaviour of the wind (e.g. different upstream terrain roughness may induce differences in mean wind speed at hub height). Additionally, the turbulence that the turbine and met mast are exposed to will not be the same; for example, the turbine could be exposed to a local gust causing the power production to increase for a period of time while the recorded wind speed is unchanged.

The other wind directions will have some combination of these two effects illustrated in Figure 3.

All of these effects, including the ones described in this section and the previous two (i.e. 2.2.1 and 2.2.2) decrease the correlation between wind speed and power and place data points above and below the fixed point power curve. Decreasing the correlation between these two parameters increases the uncertainty in the estimation of the fixed point power curve and makes it more difficult to estimate.

2.3 Existing Method for Power Curve Measurements

The current industry standard method is described by IEC 61400-12-1 [1]. This method is simple and widely used, however it has some problems [6]. A new edition of the standard is currently under review which attempts to correct some of the problems with the standard method. The method and its potential errors will be discussed in this section.

The current version of the IEC method utilizes 10 minute means of wind speed at hub height and electrical power production. The wind speed and power measurements are first filtered to remove periods where the turbine is not operating normally, and periods where the met mast is in the wake of any surrounding wind turbine. The measurements are then corrected based on the air density calculated using temperature and pressure of the air. The 10 minute data points of wind speed and power are then grouped into “bins” based on the wind speed measurements. The wind speed bins have a width of 0.5 m/s. The corrected 10 minute average wind speed and power measurements for a given bin are then averaged to obtain the representative wind speed and power for that bin. The average wind speed and power for each bin can then be plotted together and form an estimate of the steady state power curve.

The effects that the three difficulties associated with power curve measurements, described in Section 2.2, in relation to this IEC method, will be described in the Sections 2.3.1, 2.3.2, and 2.3.3.

2.3.1 Systematic Error in IEC Method Due to Turbine Dynamics and Power Curve Non-linearity

Since the driving wind speed is continuously fluctuating, the wind turbine power production is also fluctuating as described in Sections 2.2.1. This IEC method relies on the assumption that the fluctuations in wind speed and power will be symmetric about the fixed point power curve. However, there are several plausible situations where the fluctuations will not be symmetric causing the mean to deviate from the fixed point power curve.

The time series of instantaneous power production and wind speed are denoted by p and u, respectively. In the IEC standard, the 10 minute mean power, $\langle p \rangle$, is assumed to correspond to the fixed point power for the 10 minute mean wind speed, $P_{FP}(\langle u \rangle) = \langle p \rangle$. However, due to the non-linearity of the power curve, this assumption is not always valid, $P_{FP}(\langle u \rangle) \neq \langle p \rangle$ [7]. In the non-linear regions of the power curve, the mean will deviate from the fixed point power curve. Regions with positive curvature will be overestimated (typically the region at low wind speeds below rated speed), while regions with negative curvature will be underestimated (typically around the rated wind speed) as illustrated in Figure 4.

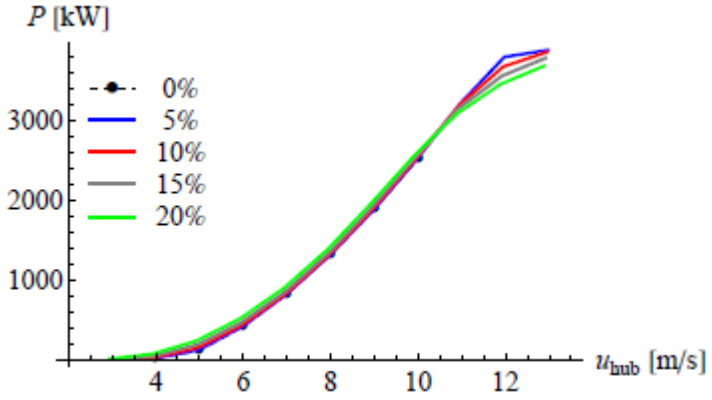

Figure 4 – Typical systematic error in power curve generated using IEC method caused by turbulence at different levels of turbulence intensity [6]

Attempts have been made to try to correct for this non-linearity using a Taylor Expansion of the fixed point power curve [7] as shown in Equation (2).

$$P_{FP}(\langle u \rangle) = \langle p \rangle - \frac{\sigma_u^2}{2}\left[\frac{\partial^2}{\partial V^2} P_{FP}(V)\right]_{V=\langle u \rangle} \tag{2}$$

However, this approach assumes weak turbulence intensities which are not typically seen in atmospheric flows. Additionally, to obtain the second derivative of the power curve, the power curve must already be known so some iterative approach with an initial guess of the power curve shape is required.

In addition to non-linearity of the power curve, the mean power production could deviate from the fixed point power curve due to differences in the response of the wind turbine between gusts and lulls in the wind speed. It is possible that the turbine could respond faster to a gust reaching steady state sooner than it would for a drop in wind speed of the same magnitude. This example would cause the mean power production to overestimate the fixed points of the power curve.

Using the method described in the current IEC standard, these effects are not considered. Therefore, to reduce the uncertainty induced by large fluctuations of the driving wind speed, the 10 minute time periods which were "too turbulent" are sometimes left out. This requires a selection of a cut-off turbulence intensity; however any level of turbulence could cause the mean power production to deviate from the fixed point power curve. This makes it difficult when choosing an appropriate cut-off turbulence level where the mean power production will not deviate "too much" and requires some engineering judgment.

A more systematic approach of normalizing the power curve to a specific turbulence level [8] has been implemented in the next revision of the IEC standard. However, this method must make prior assumptions on the shape of the power curve. Additionally, the method must make the assumption that the wind turbine is always operating on the fixed point power curve and has no method of dealing with the time lag between a change in wind speed and the turbine reaching a new steady state. If a power curve has an abnormal shape or is slow to respond to fluctuations in wind speed, the method may induce further error to the resulting power curve.

2.3.2 Error Due to Neglect of Wind Shear

The current version of the IEC standard has no way of dealing with the problems described in Section 2.2.2 related to wind shear. Wind speed measurements used for power curve validation are performed using a single cup anemometer at the wind turbine's hub height and it is assumed that this measurement is representative of the driving wind speed. However, using this assumption, it is possible that during two different measurement periods, one with a low wind shear and one with extreme wind shear, the hub height wind speed could be the same but the driving wind speed would be significantly different. An example comparison between normal wind shear ($\alpha = 0.2$) and an extreme wind shear ($\alpha = 0.8$) with the same hub height wind speed of 10 m/s is shown in Figure 5.

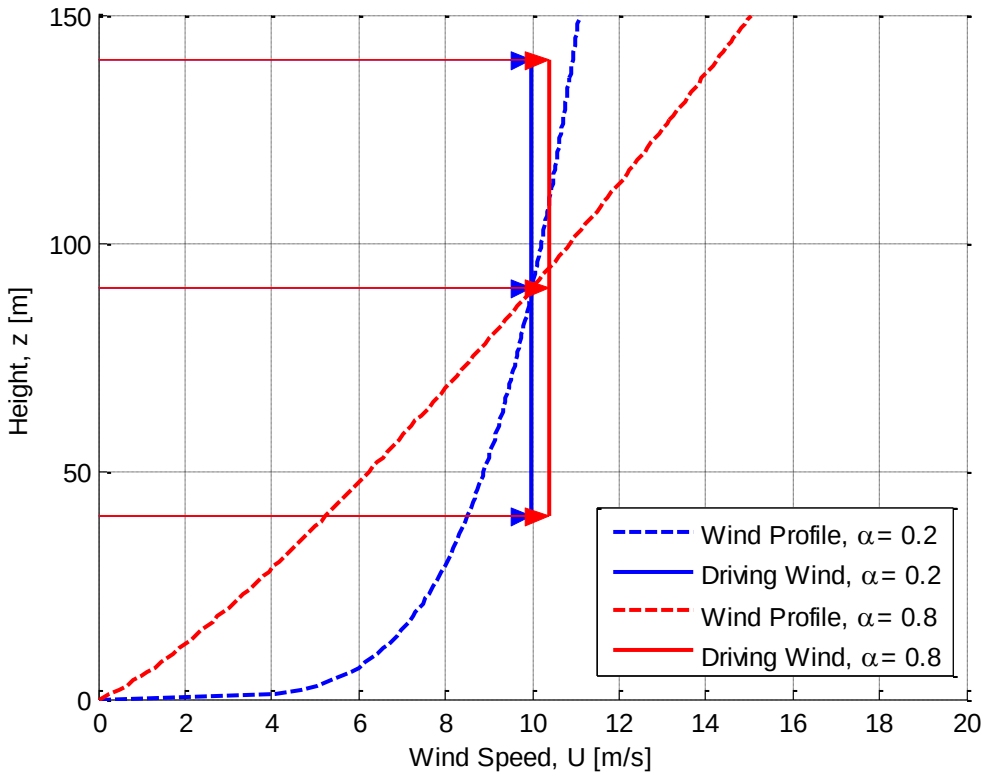

Figure 5 – Comparison of equivalent driving wind under normal and extreme wind shears with the same hub height wind speed

As illustrated by Figure 5, variations in wind shear can significantly affect the driving wind speed, and, therefore, the power production of a wind turbine. The driving wind speeds of the two wind speed profiles differ by 0.4 m/s but this difference would not be captured using the current IEC standard since only the hub height wind speed is measured and considered. This allows the power curve to vary depending on how much wind shear is at a particular site.

In the next revision of the IEC standard, a method for normalizing the wind speed has been added to help reduce the power curve's sensitivity to wind shear. A method for estimating the driving wind speed was developed [6] by averaging the wind speeds measured at several heights across a wind turbine's rotor area. One of the methods calculates, as shown in Equation (3), the average kinetic energy flux of the wind through the rotor and calculates the "equivalent" wind speed, U_{eq}, which would produce that kinetic energy flux.

$$U_{eq} = \sqrt[3]{\frac{1}{A}\int_A U(z)^3 \cdot dA} \tag{3}$$

The 10 minute mean wind speeds at several heights can then be used providing some spatial averaging. The effects of differential wind speed over the rotor area are normalized out providing a better approximation of the driving wind speed over the 10 minute period.

2.3.3 Insensitivity to Separation of Wind Turbine and Met Mast

The IEC method does not deal well with the difficulties related to fluctuating driving wind speed and uneven inflow as described above. However, the IEC method does effectively reduce the problems associated with the spatial separation between the met mast and the turbine described in Section 2.2.3. Assuming the terrain is homogeneous in all directions, the entire area will be subject to similar driving mesoscale forces resulting in similar mean wind speed and turbulence intensities over the whole area. The turbulence is continuously evolving and the turbulence is varying perpendicular to the wind direction, as discussed in Section 2.2.3. Due to these two phenomena, the two points will not be subjected to the exact same turbulence structures (e.g. eddies and gusts); however, the two points will be subjected to turbulence of the same characteristics. Over time, they will average out to the same mean wind speed and turbulence intensity. This means the IEC method is not sensitive to the wind direction and the separation between the met mast and wind turbine since it does not require second to second correlation of the wind speed and power measurements.

2.4 Proposed Method for Power Curve Measurements

A new method for estimating a wind turbine's fixed point power curve has been developed over the past several years; predominantly by a group from the University of Oldenburg. This method utilizes the high frequency data along with some more sophisticated statistical techniques to try to better estimate the power curve. The aim of this method is to improve the speed and robustness of the estimation of the steady state power curve over the standard IEC method described in Section 2.3. The implementation, practical application, and performance of this method are the main topics of investigation for this thesis.

The method assumes that the wind turbine's power output can be modeled using a stochastic differential equation called the Langevin equation. This equation models the fluctuations in power production of a wind turbine as two parts; a deterministic part driven by the wind speed and a stochastic component causing random fluctuations in the power output. The measurement data is used to fit the modeling coefficients to the data. Once the model has been fitted to the data, the fixed points of the power curve can be estimated. The implementation of this method will be described in much greater detail in Section 4.

The benefits and applications of this method are described in Sections 2.4.1 and 2.4.2, respectively.

2.4.1 Potential Benefits of Proposed Method

The Langevin method claims several advantages over the standard IEC method which are summarized as follows.

***Better Utilization of Available Data**:* Modern wind turbines are typically sampling data at a high frequency; however, this data is not effectively used since it is often immediately reduced to 10 minute averages. For example, when using the IEC method, a 10 minute data set is reduced to one point. This throws away potentially useful information contained within the time series. Additionally, each 10 minute average is treated as an independent data point. The high frequency measurements are not independent since they are connected to one another in time. Information about the dynamics of the power production is contained with the relationship in time between the points in the time series. This "Langevin method" hopes to take advantage of the information contained within the 10 minute time series to achieve a faster and more robust estimate of the power curve.

***Robustness to Turbulence**:* This Langevin method has the ability to filter out the noise using the stochastic component of the equation and extract the underlying dynamics described by the deterministic part of the equation. This provides an effective method of dealing with the difficulties described in Section 2.2.1 associated with the fluctuations in the driving wind speed. In theory this allows the method to use high turbulence data and obtain the same fixed point power curve as if low turbulence data were used. Therefore, no corrections will be required to normalize the power curve based on the turbulence intensity as needed in the IEC method [9].

Correction for Wind Shear: The Langevin method can also be corrected for wind shear as done for the IEC method described in Section 2.3.2. The instantaneous wind speed measurements at several different heights can be averaged over the rotor area of the wind turbine to estimate the driving wind speed.

Information about Wind Turbine Dynamics: After the Langevin model has been fit to the data set, the predicted fixed points are extracted from the model to predict the power curve. After this, the model can be used further to predict the response of the wind turbine's power output to a dynamic fluctuating driving wind speed [10].

Since this method has ways of dealing with the problems described in Section 2.2, the method's output is expected to be more robust to turbulence and site specific effects [9] providing a more certain power curve. Additionally, since the data is more effectively used, the fixed point power curve is expected to converge sooner and require less data than the standard IEC method. Investigation of these claimed benefits was one of the main investigations of this thesis described in Section 6.

2.4.2 Potential Applications of Proposed Method

If this Langevin method can live up to its claims, some of the potential uses for the method are summarized below.

Faster Testing and Optimization of Prototype Wind Turbine Components or Controllers: Initially, the method will have to prove itself reliable before it can be accepted by the wind energy industry. For industry to try it out and become comfortable with the method, it could be used as a method to quickly determine if a component added to a wind turbine or a change in the control strategy has positively or negatively affected the power curve.

More Effective Condition Monitoring of Wind Turbines: A wind turbine's power curve is monitored over time and compared to other wind turbines within the same wind farm. If one of the power curves is deviating from its historic performance or from the other similar wind turbines in the wind farm, it is an indication that something may be malfunctioning or wearing

out on the turbine. If less data is required to identify (with statistical significance) a change in a wind turbine's power curve, potential maintenance issues can be identified sooner and resolved [9].

More Accurate Power Forecasting: A lot of effort goes into removing the effects of turbulence during power curve validation to obtain the fixed point or "zero-turbulence" power curve. However, it should be remembered what the purpose of calculating and validating a power curve is in the first place. As explained in Section 2.1, a power curve is necessary for wind farm operators to determine how much power an existing or future wind farm will produce. The turbulence affects the amount of power produced by a wind turbine and therefore to accurately predict AEP or next day power production from a wind farm it should be included. However, the turbulence at the wind turbine's certification site will likely be significantly different from the turbulence at a potential or existing wind farm site. Therefore, using the IEC method, the effects of the certification site's turbulence must be subtracted from the power curve and then the effects of the wind farm site's turbulence must be added back in. Using the Langevin method provides a simpler solution since, to obtain the fixed point power curve, a model for predicting the dynamics of the wind turbine's response to a fluctuating wind speed is generated in the process. This Langevin model can be applied to simulated wind speed time series at the turbulence level of the site in question to generate more certain predictions of future power production [10].

Faster Verification of Wind Turbine Power Curves: Eventually, after the method has become established and accepted in the wind energy industry, it could be used as an alternative to the IEC standard method, providing a faster and more robust method of validating the power curve of prototype wind turbines during certification.

All of the benefits described above can be achieved without significant costs since the required data is already being recorded. The new method of analyzing the data would have to be added to a wind turbine's SCADA system. More data storage may be required, however, especially with the continuous reduction in the size and cost of computer memory, the benefits of implementing this method would likely outweigh this small cost.

3 Experimental Data

This section describes the experimental data used for all of the analyses described in Sections 5 and 6. A description of how the raw data was measured was described in Section 3.1. This raw data was then filtered and corrected as described in Sections 3.2, 3.3, and 3.4. Unless otherwise stated, the resulting data was the base data set which all of the analyses started with. Any additional filtering and pre-processing that was performed for each of the test cases is described in the analysis sections below.

3.1 Measurement Setup

The data used for the analysis in this thesis was taken from the 500 kW Nordtank turbine owned by DTU located at the Risø campus. For simplicity, in this report the turbine will be referred to as "the Nordtank turbine". The turbine has been the subject of many different research projects over the years and is very well instrumented and understood. Additionally, there is a well instrumented met mast located approximately 2.5 rotor diameters upstream of the turbine in the prevailing wind direction. The layout of the site including the Nordtank turbine and the met mast is shown in Figure 6. The Nordtank turbine is the largest in a row of three wind turbines which are also indicated on the figure.

Figure 6 – Risø turbine row and met mast layout

As shown in Figure 6, there is a met mast approximately 100 m from the turbine at an angle of approximately 283°. The Nordtank turbine and met mast are shown in Figure 7 (a) and (b), respectively.

Figure 7 – (a) Nordtank turbine and (b) met mast at Risø site taken from Nordtank nacelle looking towards the North West

The Nordtank turbine shown in Figure 7 (a) is a 500 kW fixed speed, stall regulated wind turbine. It had a rotor diameter of 42 m with a hub height of 36 m. It has a relatively limited control system since it does not have a pitch system or a power converter.

The met mast to the west of the Nordtank turbine shown in Figure 7 (b) is very well instrumented as summarized in Table 1.

Table 1 – Summary of met mast instrumentation

Height	Instrumentation	Note
16.5 m	Sonic anemometer on a South boom	Lower Tip Height
18 m	Two cup anemometers on North and South booms	
27 m	Cup anemometer on a South boom	
34.5 m	Sonic anemometer on South boom	
36 m	Two cup anemometers on North and South booms	Hub Height
45 m	Cup anemometer on a South boom	
52.5 m	Sonic anemometer on a South boom	
54 m	Two cup anemometers on North and South booms	
57 m	Cup anemometer centered on top of mast	Upper Tip Height

Three months of high frequency measurements were available recorded from October to December of 2011. The wind turbine measurements and met mast measurements were recorded at 35 Hz.

Since a measurement frequency of 1 Hz was recommended for use with the Langevin method, the data was immediately reduced to 1 Hz by taking 1 second averages of the raw 35 Hz data. This helps greatly speed up analysis by reducing the number of data points required for the analysis.

3.2 Filtering

The data was filtered based on several criteria to attempt to remove the measurements when the turbine was not operating normally or when the instrumentation was not measuring accurately. Some of the filtering was done based on the 10 minute means, and some was done directly on the instantaneous data.

The data was filtered based on some of the wind turbine's status signals. The status signals for the tip brakes, mechanical brake, and generator were used to ensure that neither of the brakes were applied and the generator was operating normally. The instantaneous values of the status signals were used allowing for part of a 10 minute time series to be included if, for example, the generator shut off due to low wind speeds part way through the 10 minute period.

The turbine was also filtered based on the wind direction. The sonic anemometer just below hub height (at 34.5 m) was used for the wind direction measurement. To choose the acceptable sector, the difference between the North and South 10 minute mean wind speed measurements taken at the turbine's hub height of 36 m were plotted against the 10 minute mean wind direction as shown in Figure 8

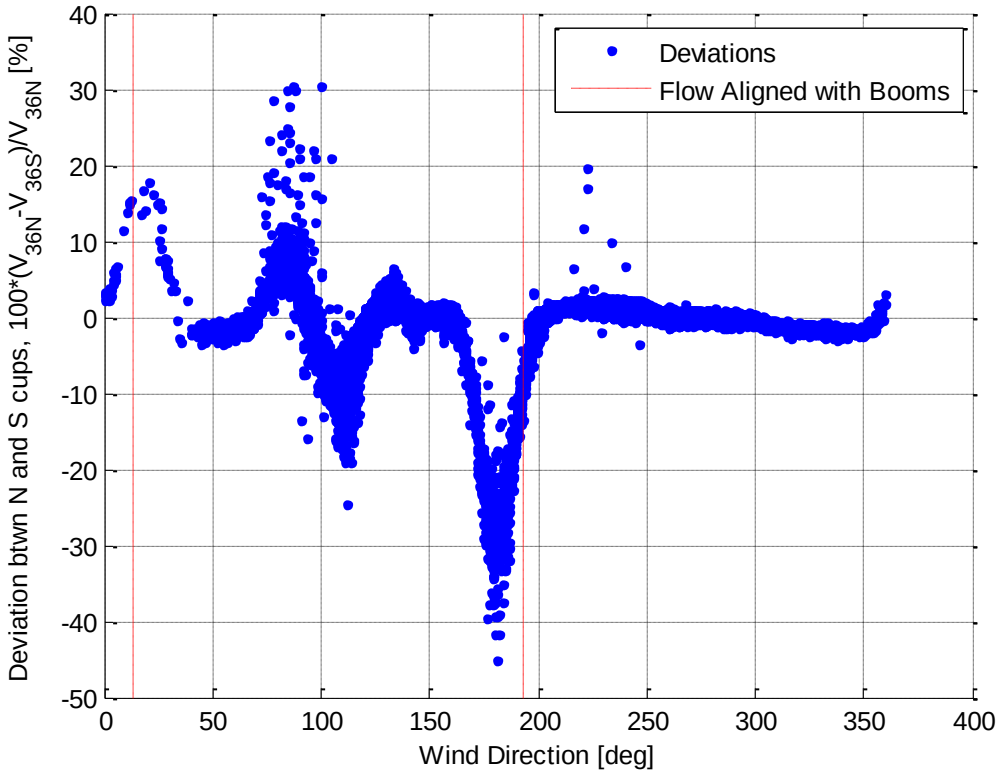

Figure 8 – Deviation between north and south cup anemometers placed at 36 m on Nordtank met mast at different wind directions

As seen in Figure 8, the two measurements differ significantly at several different wind directions. At approximately 13° and 193° degrees, the wind speed is aligned with the booms carrying the cup anemometers. This means that one of the two cups will be in the wake of the mast causing significant measurement error. Furthermore, when the wind is coming from the east (between 0°

and 180°), the met mast can be affected by the wake of one or more of the turbines in the turbine row. Due to these two effects measurements with 10 minute mean wind directions near alignment with the booms or from the east were excluded since the wind speed measurement was likely to contain significant error. All measurements with 10 minute mean directions outside the region from 210° and 360° were excluded. Also, typically additional directional filtering was performed to reduce the sector even further. The sector used for each of the analyses is specified in each of the report sections below.

Additionally, the measurements were filtered based on the 10 minute average values from the cup anemometers to ensure the wind speed was between 2 m/s and 100 m/s. Outside these limits it was assumed that the cup anemometer was not functioning correctly. These wind speeds are outside the range between the cut-in and cut-out wind speeds so the turbine should not be operating in this range anyways.

Additional filtering was done as needed for each of the test cases described in the sections below.

3.3 Equivalent Wind Speed Correction

As described in Section 2.2.2, the wind turbine is effectively averaging the wind speed over the rotor area of the wind turbine. Therefore, the effects of wind shear can significantly alter estimations of fixed point power curves adding uncertainty. Since this mast is so well instrumented, in an attempt to remove the effects of wind shear, average wind speed over the rotor area was used to better approximate the driving wind than simply using one of the measurements. The wind speed was averaged to calculate the equivalent wind speed providing the same kinetic energy flux through the rotor disc.

The method used was taken from previous work characterizing the effects of wind shear on the power curve [6]; however, the method was applied to instantaneous wind speed measurements rather than 10 minute average values. All of the cup anemometers on the met mast were used to approximate the equivalent wind speed. The sonic anemometers were not used for wind speed estimates to prevent mixing types of instrumentation. Each cup anemometer was assigned an area of the rotor. The sectors were assigned to each cup based on the cup's height and the sectors were divided at the mid-point between two cups as shown in Figure 9.

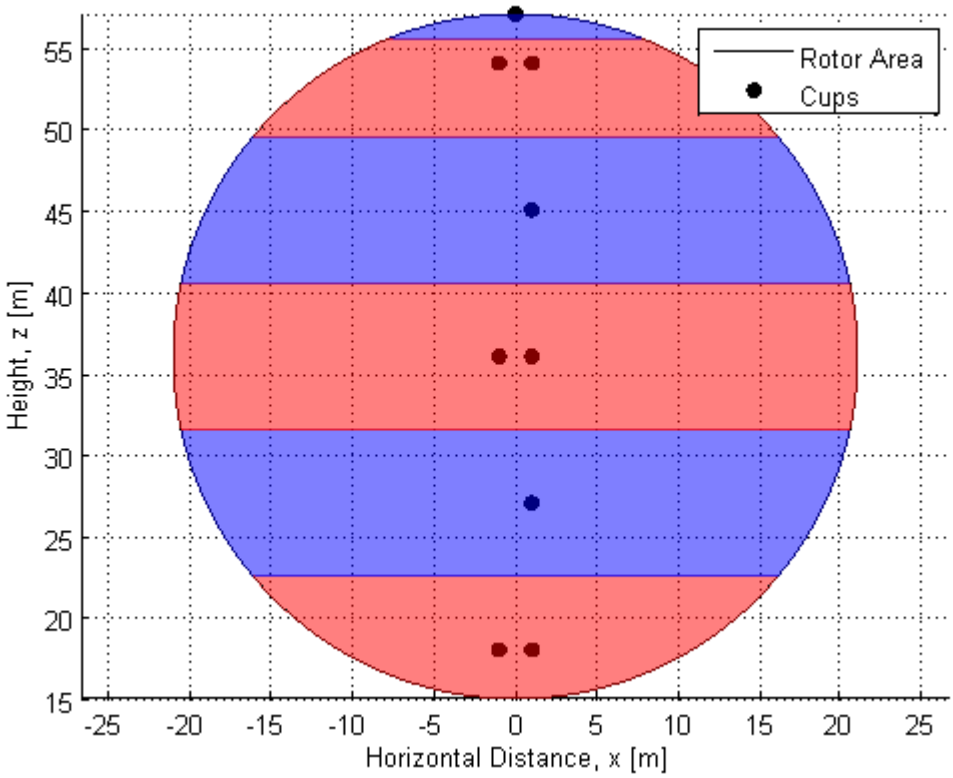

Figure 9 – Division of rotor area assigned to each cup

The wind speed measurements were assumed to be constant across their assigned area. For areas with two cups, the average of the two cups was used. The equivalent wind speed was then calculated as shown in Equation (4).

$$u_{eq}(t) = \sqrt[3]{\frac{1}{A}\sum_{n=1}^{N_{sectors}} A_n \cdot [u_n(t)]^3} \tag{4}$$

This equivalent wind speed was used in all the analyses presented below unless stated otherwise.

3.4 Air Density Correction

To remove the effects of changes in atmospheric pressure and temperature, the air density could be corrected to standard temperature and pressure. However, reliable temperature and pressure measurements were not available for the period in question so the correction was not performed. This is not expected to significantly impact the results since the turbine is nearly at sea level and the temperature variation over the period is not expected to vary greatly. In previous work, the density correction was also neglected and the Langevin method was still implemented with success [3]. This assumption should be checked in future work as described in the recommendations described in Section 8.

4 Langevin Reconstruction of Power Production Dynamics for Power Curve Validation

The Langevin equation is a stochastic differential equation which has been used for many things in the past. It was originally developed by Paul Langevin in 1908 to describe Brownian motion of a particle in a fluid. Recently it has been applied to modeling wind turbine power production and a method for extracting a power curve from high frequency measurements of wind speed and power, predominantly at the University of Oldenburg [2] & [3]. This method developed at the University of Oldenburg has been implemented as described in this section and then further developed and tested as discussed in the following sections.

The Langevin method is effective at extracting underlying trends from noisy fluctuating time series. The wind speed measurements required for power curve validation are continuously fluctuating due to fluctuation in the driving wind speed. Additionally, the measurements can be considered to be noisy, but not in the traditional sense from electrical interference or problems with the instrumentation. As described in Section 2.2, differences between the driving wind speed at the rotor and the measured wind speed can be caused by turbulence, wind shear, and the spatial separation between the measurement point and the wind turbine. This difference can be thought of as measurement noise on top of the true driving wind speed which we are trying to measure.

This section aims to explain how this model works as well as how it was implemented for this report. The source code for the implementation of this Langevin method has been made available for future use in the FastWind project. This section aims to help understand the methodology and required inputs so the source code can be used or improved for later stages of the FastWind project. The model which describes the power production of a wind turbine using the Langevin equation is described in Section 4.1. The method used to fit this model to a set of measurement data is described in Section 4.2.

Finally, to see if the Langevin method has been correctly implemented, the algorithm was tested using some simulated data as described in Section 4.4. Simulated data was used since the theoretical solution is known as it must be chosen by the user when generating the simulated data.

4.1 Langevin Model for Wind Turbine Power Output

The Langevin method works by fitting a stochastic model to the wind speed and power measurements. This model uses the Langevin equation to split the fluctuations of power output of the wind turbine into a deterministic component and a stochastic noise component. The deterministic component is always moving the power production of the wind turbine towards its steady state value for the current driving wind speed. The stochastic noise component is offsetting the power production from the steady state value causing randomness in the fluctuations.

The one dimensional Langevin stochastic differential equation used in the model is shown in Equation (5). This equation relates the rate of change of the power, $\dot{p}$, at a given time, t, with the current power production, p.

$$\dot{p}(t) = D^{(1)}(p) + \sqrt{D^{(2)}(p)} \cdot \Gamma(t) \tag{5}$$

The first term, $D^{(1)}$, is referred to as the drift coefficient and is a function of the power production. The drift term is the deterministic part of the equation which drives the power production back towards steady state. The second term contains two components, $D^{(2)}$ and Γ, which are referred to as the diffusion coefficient and Langevin force, respectively. The Langevin force is the stochastic component with a mean of 0. The diffusion coefficient is also a function of the power production and dictates how much the Langevin force influences the power production.

The drift and diffusion coefficients are referred to as coefficients, implying that they are constants; however, they are more like functions since they depend on the instantaneous power production. To illustrate the role of the two coefficients a set of simple fictional drift and diffusion equations have been created as shown in Equations (6) and (7), respectively.

$$D^{(1)}(p) = 1 - p \tag{6}$$

$$D^{(2)}(p) = 1 + p^2 \tag{7}$$

The resulting drift and diffusion curves were plotted as shown in Figure 10.

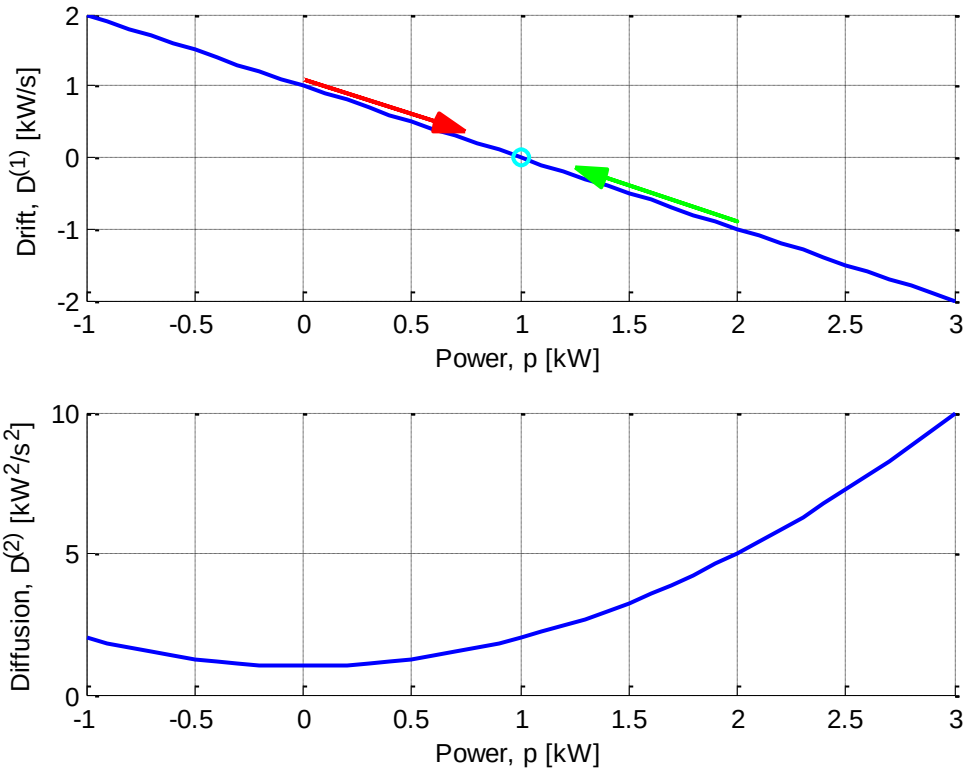

Figure 10 – Fictional drift and diffusion curves

In the region where the drift curve shown in Figure 10 is positive (i.e. $p < 1$), the rate of change of power, $\dot{p}$ calculated using Equation (5), will also be positive on average since the Langevin force, Γ, in the diffusion term has a mean of zero. On the contrary, in the region where the drift curve is negative (i.e. $p > 1$), the rate of change of power will be negative. This means that the power production will, on average, always be moving towards the zero crossing in the drift curve (in this case $p = 1$). This zero crossing represents the fixed point power for this drift curve.

To illustrate this restoring force pushing the power production towards the fixed point value, two simulated time series were generated starting below and above the fixed point at p of -1 and 3, respectively. The resulting time series are plotted as shown in Figure 11.

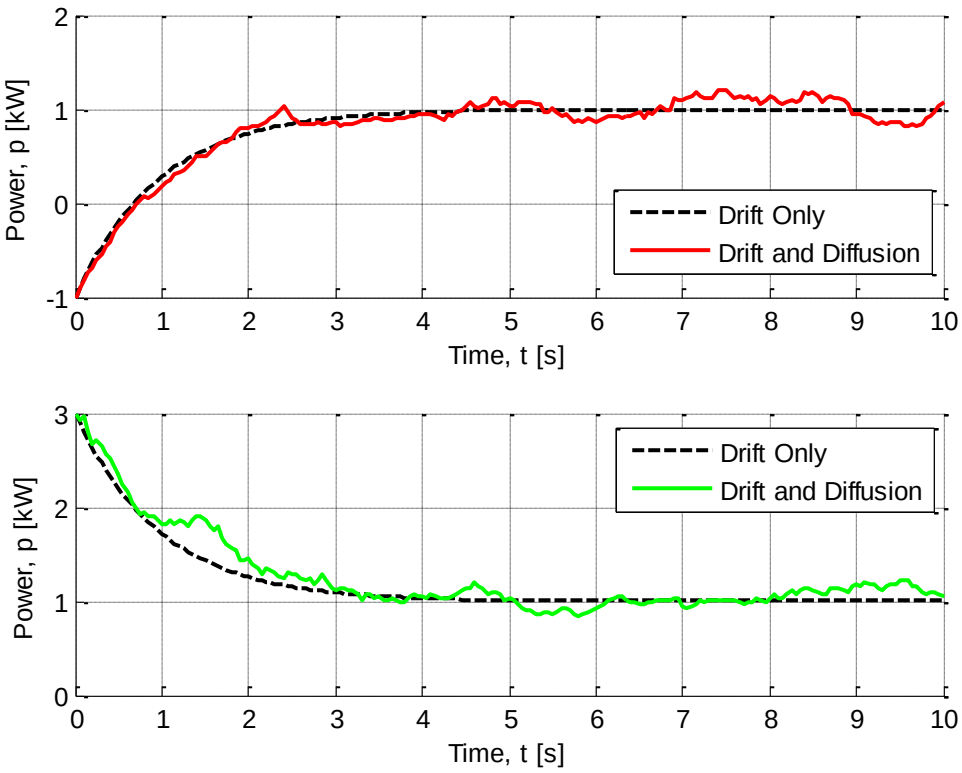

Figure 11 – Example of dynamic response for fictional drift and diffusion curves

As seen in Figure 11, the power production is forced back to the fixed point value by the drift term of the Langevin equation. The diffusion term adds random fluctuations to the time response. Looking back at Figure 10, the diffusion curve has higher values above the fixed point power than below. This means the fluctuations of the response coming down from a power production of 3 to the fixed point will typically be more severe than the response coming up from -1.

This example describing the response of the power output was performed without considering the wind speed. The drift and diffusion curves are dependent on wind speed therefore some way of incorporating the wind speed into the model is required. The wind speed and power data is binned based on the wind speed measurement. Drift and diffusion curves can then be determined for each wind speed bin to describe the response of the wind turbine's power production when the wind speed is within the bin limits. The Langevin drift and diffusion curves are independent of those in the other bins. This means that, as the wind speed fluctuates, the response of the wind turbine's power production is switching between the different Langevin curves in the model. The Langevin equation describing the i^{th} bin of the full model can be described by Equation (8).

$$\dot{p}(t) = D_i^{(1)}(p) + \sqrt{D_i^{(2)}(p)} \cdot \Gamma_i(t), \quad \text{if} \quad U_i^{(-)} < u(t) < U_i^{(+)} \tag{8}$$

This Langevin model described above must be fit to a set of wind and power data as described in Section 4.2. The drift and diffusion curves for a given wind turbine are initially unknown and must be estimated from the data. The zero crossings of the estimated drift curve for each wind speed bin can then be determined to approximate the fixed point power curve of the wind turbine.

4.2 Method of Fitting Model to Measurement Data

The Langevin model described in Section 4.1 must be fitted to wind speed and power data from a wind turbine to obtain an estimation of the power production dynamics. From the fitted model, we can estimate the wind turbine's fixed point power curve.

To perform the fitting of the Langevin model, time series data sampled at a reasonably high frequency is required; typically 1Hz has been suggested as a minimum sampling rate. Instead of treating each measurement point as an independent event, the chronological connection between the data points from the time series is used to perform the fitting to the Langevin model.

The fitting process has three main steps. The data must first be binned as described in Section 4.2.1. Next, the binned data is used to estimate the drift curve of each bin as described in Section 4.2.2. Finally, once all of the drift curves have been fit to the data, the fixed points can be estimated to obtain a power curve as described in Section 4.2.3. Once this model has been fit to the data and the fixed points have been extracted, the uncertainty of the fixed points should be estimated as described in Section 4.3.

4.2.1 Binning Data

Time series data measured with a fast sampling frequency is scattered around the zero turbulence power curve. The data point density plot for the raw data is shown in Figure 12. To perform the Langevin reconstruction, the data must be binned twice. The first binning is with respect to wind speed, since the Langevin model described in Section 4.1 has a separate drift curve for each wind speed range. The second binning is with respect to power and the reason this is necessary will be described in greater detail in Section 4.2.2.

The wind speed binning was performed with fixed width wind speed bins of 0.5 m/s as recommended in the IEC 61400-12-1 method and previous work [3]. After binning using the wind speed, a histogram of the power output data for that bin can be plotted as displayed in Figure 13. The data in this bin will be used to create a drift curve. This histogram can also be thought of like a cross section of the data density plot seen in Figure 12.

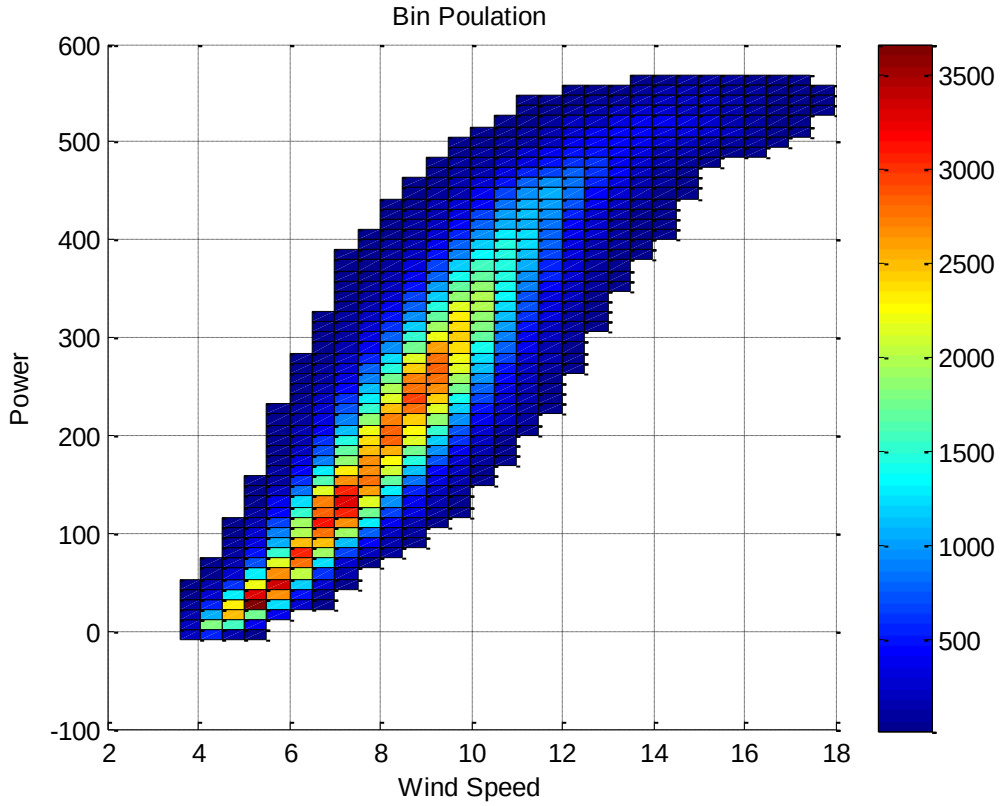

Figure 12 – Raw power curve data density plot

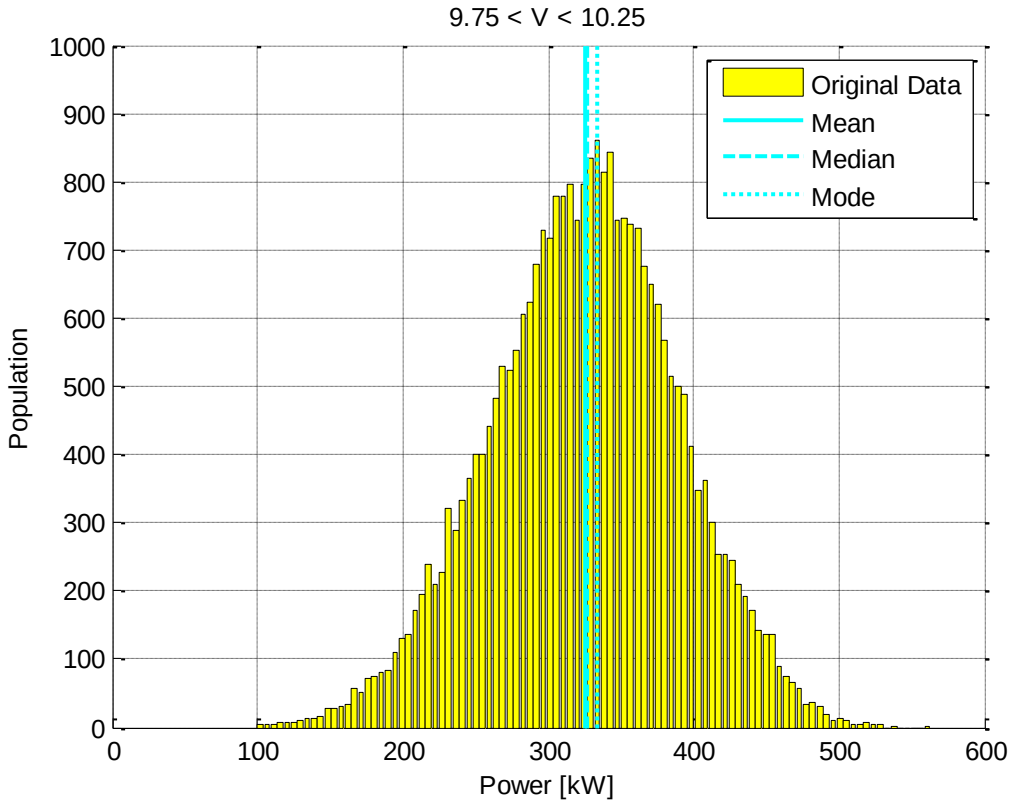

Figure 13 – Distribution and statistics of raw power data from one velocity bin

The range of wind speeds for each bin is defined, but, a representative wind speed is also required. The resulting drift curve and fixed point will be placed on this representative wind speed when plotting (e.g. for the fixed point power curve). The mean wind speed of all data contained in the bin was used as the representative wind speed. For the bin shown in Figure 13, the representative wind speed is 9.99 m/s; very close to the midpoint of 10 m/s.

Two different methods for binning in the power direction (referred to henceforth as sub-bins) were implemented to be tested. The first and simplest method is to use evenly spaced sub-bins of a chosen width for all velocity bins. In the second method, the number of sub-bins was specified and then the sub-bin sizes were adaptively selected for each velocity bin based on the data in the velocity bin. The sub-bin limits were chosen so that an even number of data points were contained in each sub-bin.

For an example, data from the Risø Nordtank turbine was binned using the two different methods for the velocity bin from 9.75 m/s to 10.25 m/s. The sub-bin populations after binning with the two methods are shown in Figure 14 and Figure 15. The widths of the blue bars show the widths of the sub-bin while the height shows each sub-bin's population.

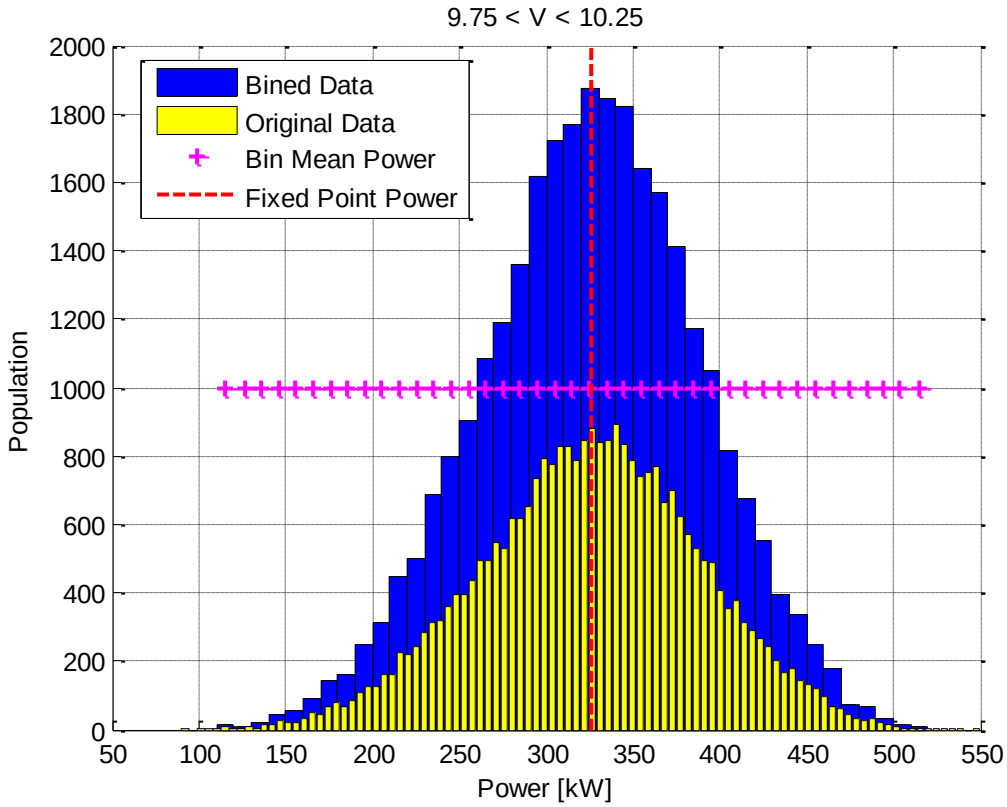

Figure 14 – Sub-bin populations using fixed width sub- bins of 10 kW

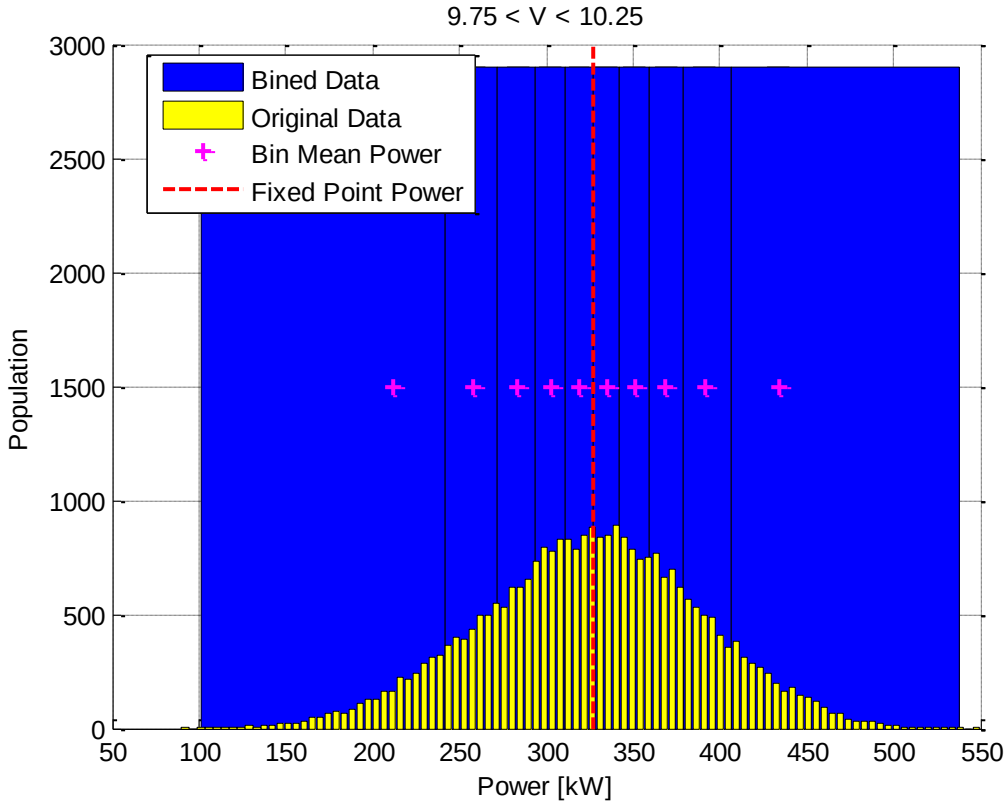

Figure 15 – Sub-bin population using adaptive bin widths split into 10 even sub-bins

As seen in Figure 14, the sub-bin populations vary greatly from over 1800 to less than 100 data points. This large variation in number of data points used to calculate the drift points causes variation of the uncertainty of the drift points within a single drift curve. This is the main weakness of binning the data in this fashion, and, for this reason, the alternative binning method shown in Figure 15 was the default method used in the analysis. The sensitivity of the fixed point power curve to the method of binning and the number of sub-bins used is described in Section 5.3.2.

Finally, similar to the representative velocity, the sub-bins require a representative power. The average power of the data in the sub-bin was used as the representative value. The representative power is shown for each bin in Figure 14 and Figure 15 as the pink markers.

4.2.2 Drift Field Reconstruction

After binning the data twice, the drift curve for each velocity bin must be estimated. However, due to the noise in the measurement data, the drift curve is difficult to determine. For this reason, the second binning with respect to power is required. This allows several points along the drift curve to be estimated with some certainty and the rest of the curve can be obtained by interpolating between the estimated points.

As explained in Section 4.1, within a single wind speed bin, the drift curve is a function of the power production. At a given power, there is a corresponding drift. Starting at this given power, the drift is the average rate of change of the power's time response. Therefore we need a way to estimate the average rate of change of power at a given power level. However, since there are significant amounts of noise in the signal, the instantaneous rate of change will be dominated by the signal's noise. For this reason, the average time response of the power production originating at a given power level is estimated. To estimate the average time response, we need a group of many time responses originating at the same initial power. This is why the second binning of the data is required, to group the time responses by their initial point.

To estimate a drift curve, one sub-bin is considered at a time. The data points included in one particular sub-bin are the starting points and the power time series after these starting points are investigated. For illustration purposes, an example sub-bin with velocity limits of 9.75 m/s to 10.25 m/s and power limits of 299 kW to 315 kW will be used. This sub-bin contains K data points and the k^{th} data point corresponds to the starting point recorded at time t_k. The time lag, τ, specifies the position in the time series after the starting point. This results in K time series segments of power which are plotted on top of one another as shown in Figure 16 (a).

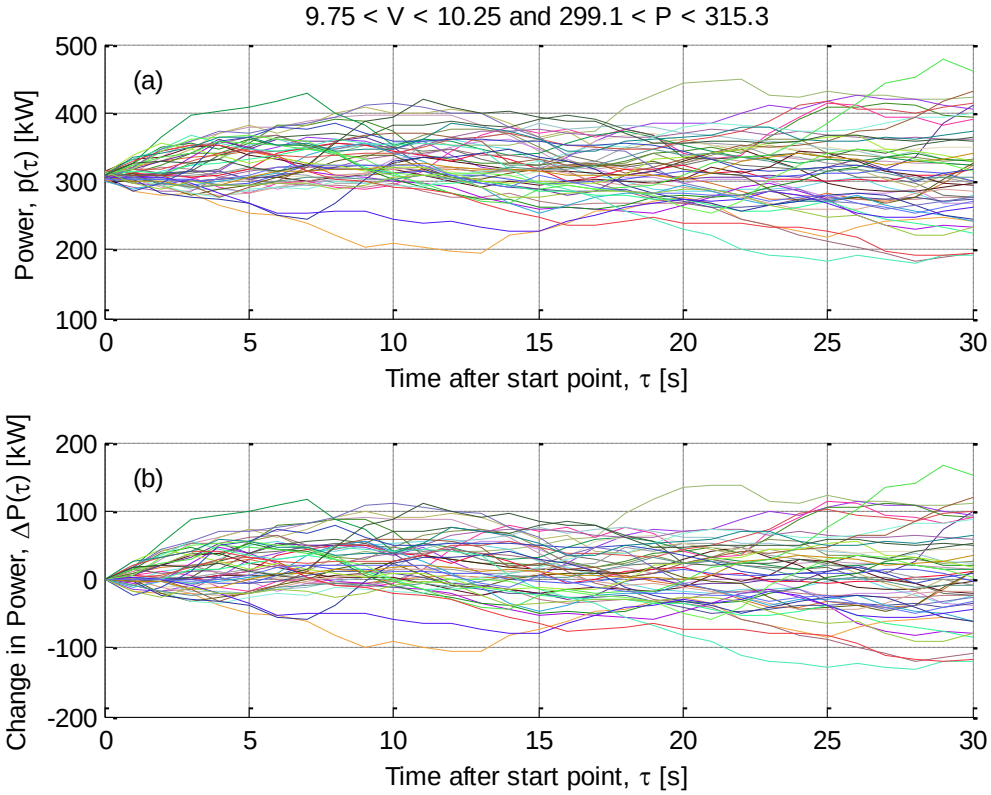

Figure 16 – Example of time series segments initiating within a chosen sub-bin

The starting points are spread out over the range of the sub bin (between 299 kW and 315 kW in the example). Therefore the change in power from the starting point is calculated as shown in Equation (9). When using the change in power, as shown in Figure 16 (b), all of the time series segments start at zero.

$$\Delta P_k(\tau) = p(t_k + \tau) - p(t_k) \tag{9}$$

The average trend of the change in power as a function of time lag is calculated by averaging the K time series segments at each τ as shown in Equation (10). This is referred to as the "averaged response" of the wind turbine under these initial conditions of wind speed and power.

$$\Delta P(\tau) = \langle \Delta P_k(\tau) \rangle = \frac{1}{K} \sum_{k=1}^{K} \Delta P_k(\tau) \tag{10}$$

This averaging results in one representative average time response for displaying how the wind turbine typically responds when the wind speed and power fall within the bin under consideration. The averaged response for the example sub-bin is shown as the black line in Figure 17.

The definition of the drift is the slope of this average response at $\tau = 0$; however, due to measurement noise and a finite sampling rate, it is difficult to accurately estimate the slope at $\tau = 0$. Instead, the slope of a best fit line for some range of tau is used as an estimate of the drift. This

introduces some uncertainty to the method as, instead of the instantaneous slope at exactly $\tau = 0$, the average slope over some range of τ is taken for the drift. The range of τ which should be used for the best fit line depends on the turbine under consideration since it depends on the speed that the turbine can respond to changes in wind speed and find a new equilibrium. It also depends on wind speed since at higher wind speeds they turbine wake dynamics are faster; however, currently, the same range of τ is used for all wind speed bins. A guideline for choosing the range of τ for fitting has not been made, and, until now, has been chosen based on experience (i.e. an educated guess). It is suggested to check some of the fits at different wind speeds to see if the range of τ chosen by the user is reasonable. This user defined selection of range of τ is not ideal; however, results can be compared for a given turbine if the same range of τ was used for both analyses. The effects of the choice of range of τ and the sensitivity of the predicted fixed points were investigated later in Section 5.3.1.

This fitting procedure for the example bin is shown in Figure 17. The range of τ used for estimating the drift was from 1s to 10s.

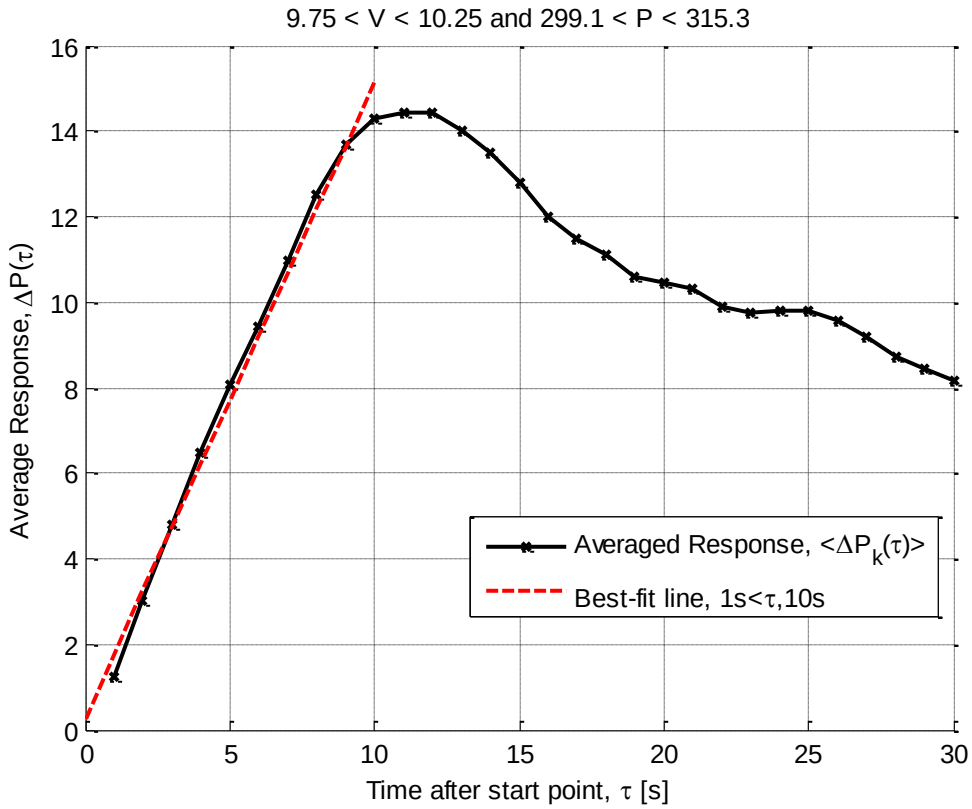

Figure 17 – Fitting of best fit line to mean changing power response to estimate drift of a sub-bin

The slope of the best fit line shown in Figure 17 is 2.77 kW/s and is the estimate of drift for this bin.

As seen in Figure 17, the response is approximately linear for the range of τ used for fitting to approximate the drift. This indicates that, for this particular bin, estimating the drift using the slope over this linear range of τ will provide a close approximation of the drift. Additionally, since the response is approximately linear for some range of τ, the drift calculated for this bin is expected to be somewhat insensitive to the choice of τ chosen provided it is still in the linear range. Finally, the range of τ used for this example bin appears to be a reasonable selection since it does not try to fit a linear function beyond 10s where the response clearly plateaus and declines.

This procedure is repeated for all of the sub-bins within the same velocity range (e.g. 9.75 m/s to 10.25 m/s in the example) to obtain drift estimates at several points along the drift curve. The estimates of the drift curve can then be plotted as shown in Figure 18. The drift curve was linearly interpolated between the drift estimations.

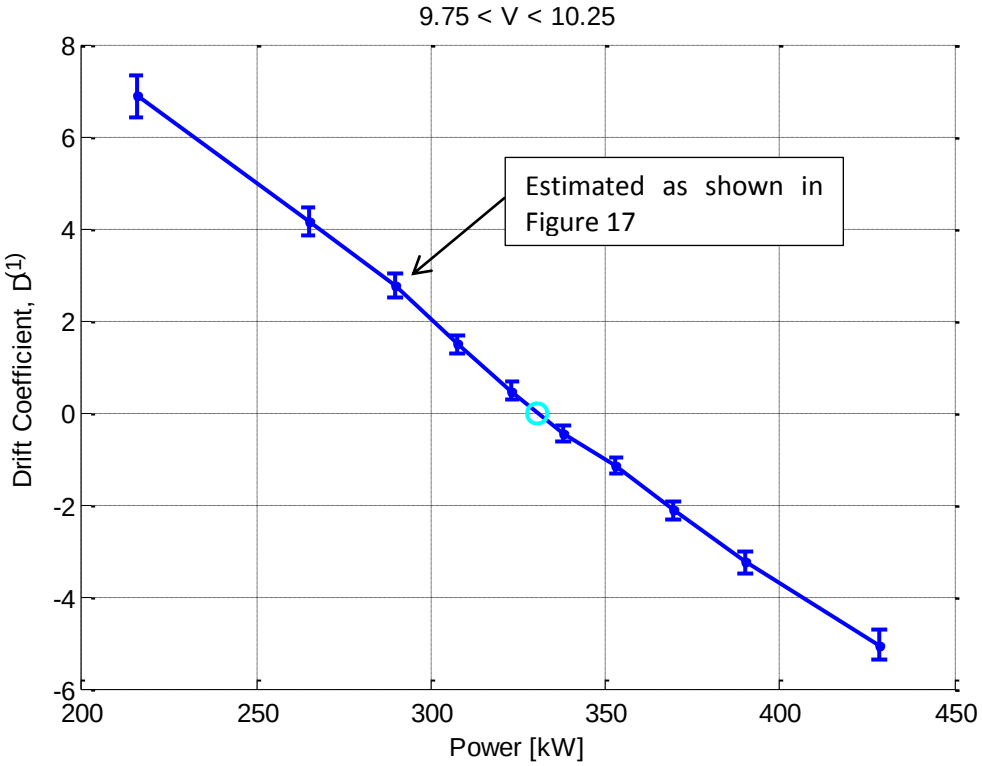

Figure 18 – Example of drift curve for a single velocity bin

Once the drift curves have been estimated for all of the wind speed bins, the fixed point power curve can be estimated as described in Section 4.2.3.

4.2.3 Fixed Points

Once the drift curves for all the bins have been determined, the fixed points can be estimated. The drift curve for each bin can be plotted as shown in Figure 18. The drift curve is calculated by fitting a linear or cubic spline to the drift points. The point where this curve crosses zero is the Langevin fixed point and can be determined by calculating the drift potential, shown in Equation (11). The fixed point for bin i, with a representative wind speed of U_i, is then the power which produces the minimum of the drift potential as shown in Equation (12).

$$\Phi_{D_i}(P) = -\int^{P} D_i^{(1)}(P^*)dP^* \tag{11}$$

$$\Phi_{D_i}(P_{FP}(U_i)) = min\{\Phi_{D_i}(P)\} \tag{12}$$

As seen in (11), the integral does not have a lower bound which makes it difficult to integrate. However, since only the relative drift potential of is needed when searching for the minimum, the lower bound is not as important. The lower bound was always set to the representative power of the lowest sub-bin of each drift curve such that the drift potential curve will always start at zero.

The drift potential for the bin between 9.75 m/s to 10.25 m/s was plotted along with the fixed point power at the minimum drift potential as shown in Figure 19.

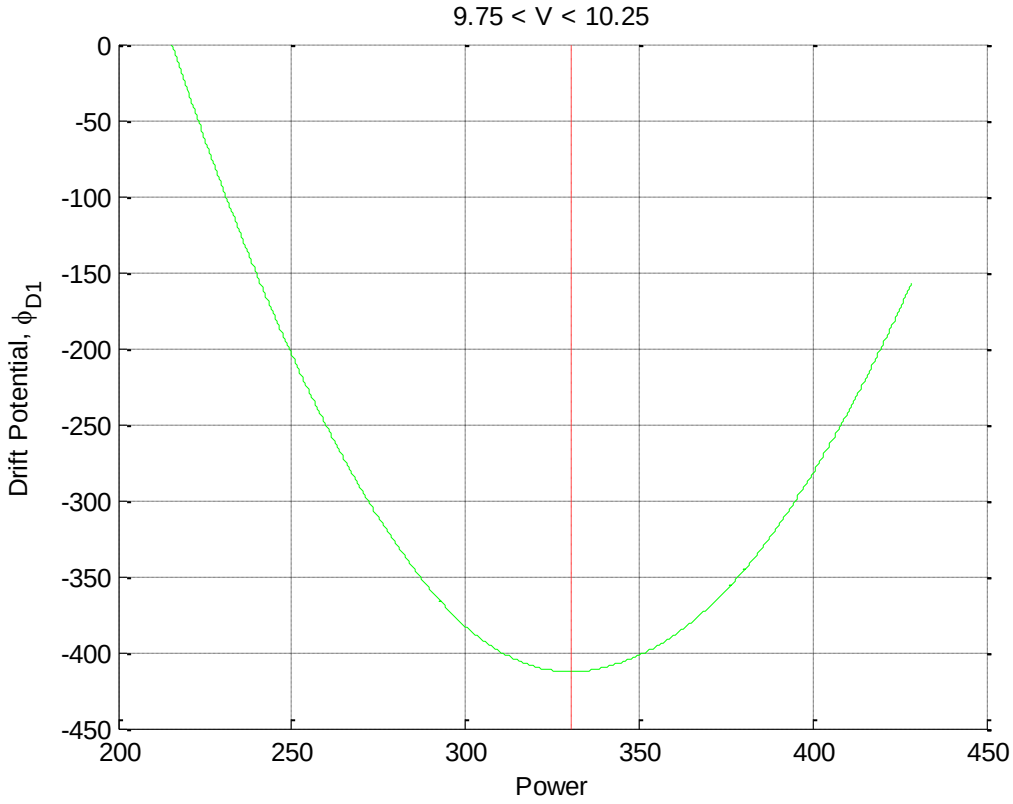

Figure 19 – Plot of drift potential for a single wind velocity bin

It was assumed that there would be only one fixed point for each bin. It is theoretically possible for a drift curve to have more than one fixed point if it has more than one zero crossing with a negative slope. However, for this simple uncontrolled turbine, no evidence was present that indicated any of the bins have multiple fixed points. Occasionally, the drift curve would fluctuate across the zero several times but it was always in the sparsely populated bins at high wind speeds and the multiple crossings were always within the range of the fixed point error bars. This just indicates that the drift points and fixed points are very uncertain in these bins since the multiple crossings are within the large range of error in the drift curve. Two distinct zero crossings in the drift curve were never observed.

The global minimum of the drift potential plots were used as the fixed points. A drift curve which has multiple zero crossings will produce several local minima in the drift potential plot. Using the global minimum provides some way to choose one as a best estimate of the fixed point.

4.2.4 Diffusion Field Reconstruction

The diffusion field was also estimated; however, the diffusion is not as important in this thesis as the drift. The diffusion field would be used when simulating the wind turbine's power output based on some measured or forecasted wind speed data [10]. Since power production forecasting was not investigated in this thesis, the diffusion coefficients were typically unused.

The diffusion coefficients are calculated using the exact same procedure as the drift coefficients with one change. Equation (10) from Section 4.2.2 would be replaced with Equation (13).

$$M^{(2)}(\tau) = \langle [\Delta P_k(\tau)]^2 \rangle = \frac{1}{K} \sum_{k=1}^{K} [\Delta P_k(\tau)]^2 \tag{13}$$

The difference being that the change in power is squared before averaging resulting in the calculation of the 2nd moment, $M^{(2)}$, of ΔP_k. A line would then be fit to the second moment as performed in Figure 17. The resulting slopes from each sub-bin would give the points of the diffusion curve.

4.3 Error Estimates for Langevin Fixed Points

After estimating the drift curves and fixed points, we need some method of determining how uncertain these fixed point estimates are. Several methods for calculating the fixed point uncertainty have been developed. One simplistic method using an analytical equation which has been used in the past is described in Section 4.3.1. Additionally, two alternative data resampling methods which attempt to estimate the uncertainty of the fixed points based on the data itself are described in Sections 4.3.2 and 4.3.3.

4.3.1 Analytic Expression Based on Langevin Parameters

An expression for calculation of the uncertainty of the drift of a given sub-bin has been used in the past [3] based on the Langevin parameters as shown in Equation (14).

$$\sigma\left[D_i^{(1)}(P)\right] = \sqrt{\frac{2}{\tau}\frac{D_i^{(2)}(P)}{N} - \frac{\left(D_i^{(1)}(P)\right)^2}{N}} \tag{14}$$

This equation relates the drift, $D_i^{(1)}$, and diffusion, $D_i^{(2)}$, coefficients to the uncertainty of the drift coefficient. The assumptions and derivation used to obtain this equation are explained in a previous paper [11]. The uncertainty of the drift points can be plotted on the drift curve as shown in Figure 20.

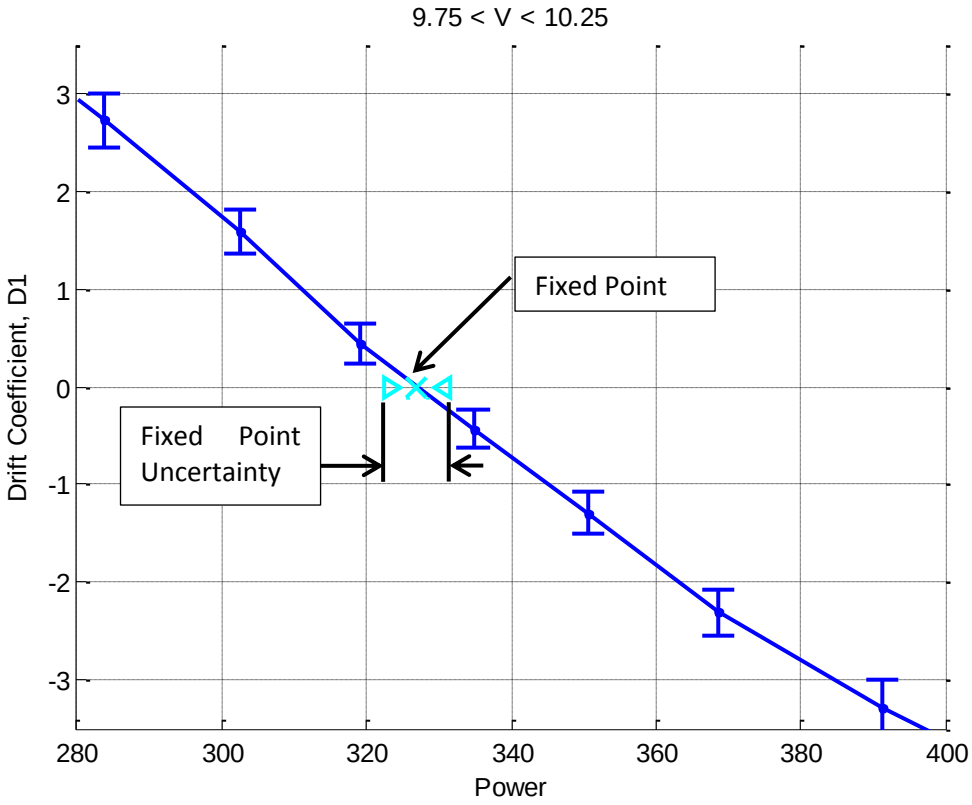

Figure 20 – Drift curve with error bars from analytic expression for bin centered on 10m/s

The error bars act as an envelope representing the upper and lower bounds of potential drift curves. This uncertainty in the drift curve can be translated into the uncertainty of the fixed point by determining the point at which the upper and lower bounds cross the zero drift point. The zero crossing of cubic or linear splines defining the upper and lower bounds of the drift curve indicate estimates for the upper and lower error bars of the fixed point.

4.3.2 The Bootstrap Power Bins Method

The above method for estimating uncertainty of the fixed points described in Section 4.3.1 assumes that the uncertainty of the drift point is related to the drift and diffusion. This can be proven for a theoretical Langevin process however, it is not possible for this real process which may not follow the Langevin model perfectly. For this reason, a method for estimating the uncertainty of the drift points from the variance of the measurement data was desired. The bootstrap power bins method uses a standard statistical resampling technique called bootstrapping [12] to estimate the uncertainty of the drift estimate from each power bin.

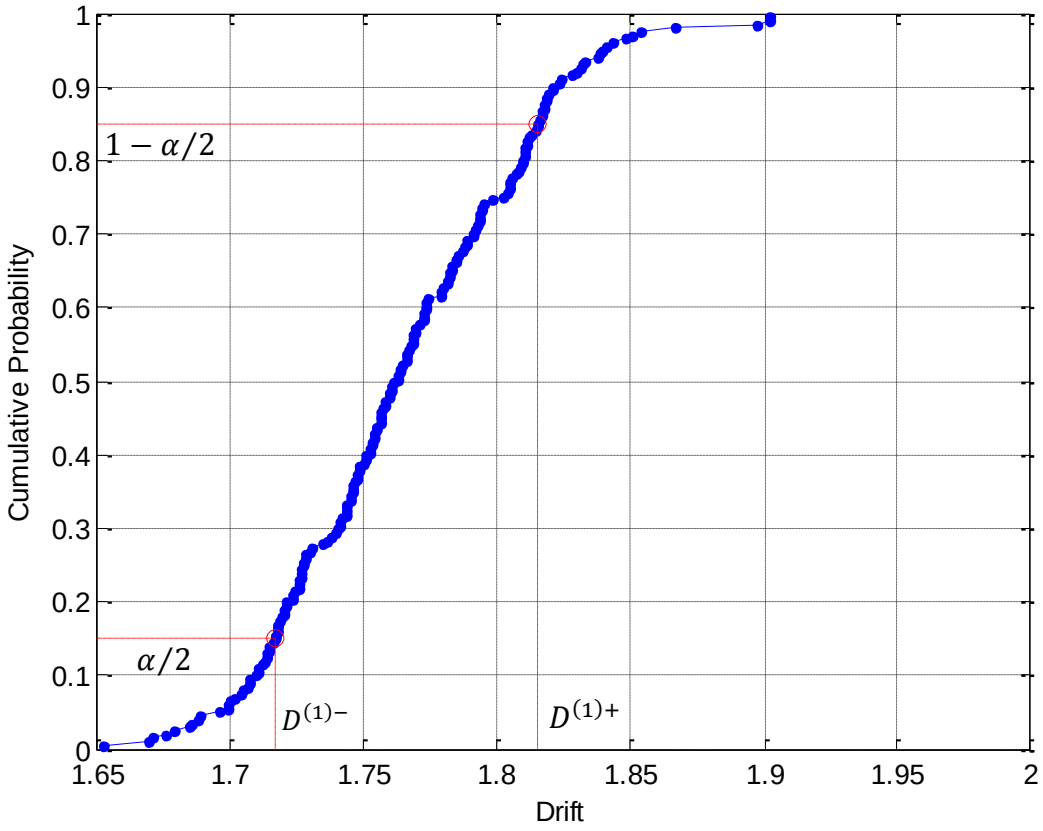

Figure 21 – Cumulative distribution function of bootstrap estimates of the drift

To perform this method, all N data points in a given bin are used calculate the best estimate of the drift. The N data points contained in the bin are then treated as a pseudo-population which is representative of the entire population of data points which were ever recorded in this bin. Simple random samples of n data points are drawn at random from the pseudo-population. This sampling is performed "with replacement" meaning that the points are replaced after they are drawn giving the possibility for them to be drawn again. This sample of n data points is referred to as a "bootstrap" sample. Using this bootstrap sample, the drift can be estimated again giving a bootstrap estimate of the drift. Many bootstrap samples are drawn from the pseudo-population providing many bootstrap estimates of the drift. These bootstrap drift estimates can then be sorted and plotted in a cumulative distribution function as shown in Figure 21.

As shown in Figure 21, a level of confidence must be chosen for the drift coefficient error bars to get the confidence interval around the drift estimate. In this example a 70% confidence interval was used corresponding to an α of 0.3.

After each bin has been bootstraped and the confidence interval around the drift point has been calculated, the envelope of the error bars can be used to convert the uncertainty of the drift curve into the uncertainty of the fixed point as was described in Section 4.3.1.

4.3.3 The Bootstrap Time Series Method

The bootstrap power bins method is good for estimating the uncertainty of the drift points for each of the sub-bins. However we are more interested in the uncertainty of fixed point estimation. The upper and lower bounds of the confidence interval around the drift points can be used to translate the uncertainty of the drift curve into the uncertainty of the fixed points as described in Section 4.3.2. However, the validity of this method of transferring the envelope of drift error into an estimate of the fixed point error has never been tested. It is possible that this method may be over- or underestimating the fixed point uncertainty.

The bootstrap power bins method does not consider the relationship in time of two points contained within a bin. It is possible that if the drift curve is changing slightly over time due to changes in the turbine or seasonal changes in wind conditions, the bootstrap power bins method would average these changes out by pulling data points from all segments of the full time series. A method which is influenced by the relationship of the points in time could more effectively capture these effects and could improve the accuracy of the uncertainty estimate.

Once again, all of the raw data is used to calculate a best estimate of the drift field and the fixed points. After this the bootstrap time series method is used to estimate the uncertainty of the fixed points. The bootstrap time series method aims at providing a more robust estimation of the uncertainty by repeating the entire Langevin drift field reconstruction for many different “bootstrap time series”. These bootstrap time series are constructed from the original data set such that they are different but have similar characteristics to the original time series.

To construct a bootstrap time series, segments of the original time series are randomly selected and extracted to be assembled into a new data set with the same length and characteristics as the original time series. The length of the segments to be extracted must be selected and, 3000 points per segment has been chosen. This segment length is estimated to be enough, plus some safety factor, for the wind speed to decorrelate. Start points within the original data set are selected at random with replacement and the following data segment is added to the bootstrap time series until the number of points in the bootstrap is as large as the original time series. The bootstrap time series is different from the original data set because some segments of the time series will be used multiple times and some segments will be omitted completely. Additionally, the number of points in each bin will vary with the bootstrap estimates and the bin limits will change since the data is re-binned for each bootstrap dataset. All of these variations give some indication of the variance of the entire fixed point estimation procedure.

The drift field is then reconstructed for each of the bootstrap time series. The fixed points for each of the drift fields are then calculated and compared. All of the estimates of the fixed points in a given wind speed bin are then compared and used to construct the probability distribution and determine a confidence interval using the same method as shown in Figure 21 from Section 4.3.2. An example of the best estimate drift curve and the bootstrap estimates along with the fixed point and confidence interval are shown in Figure 22.

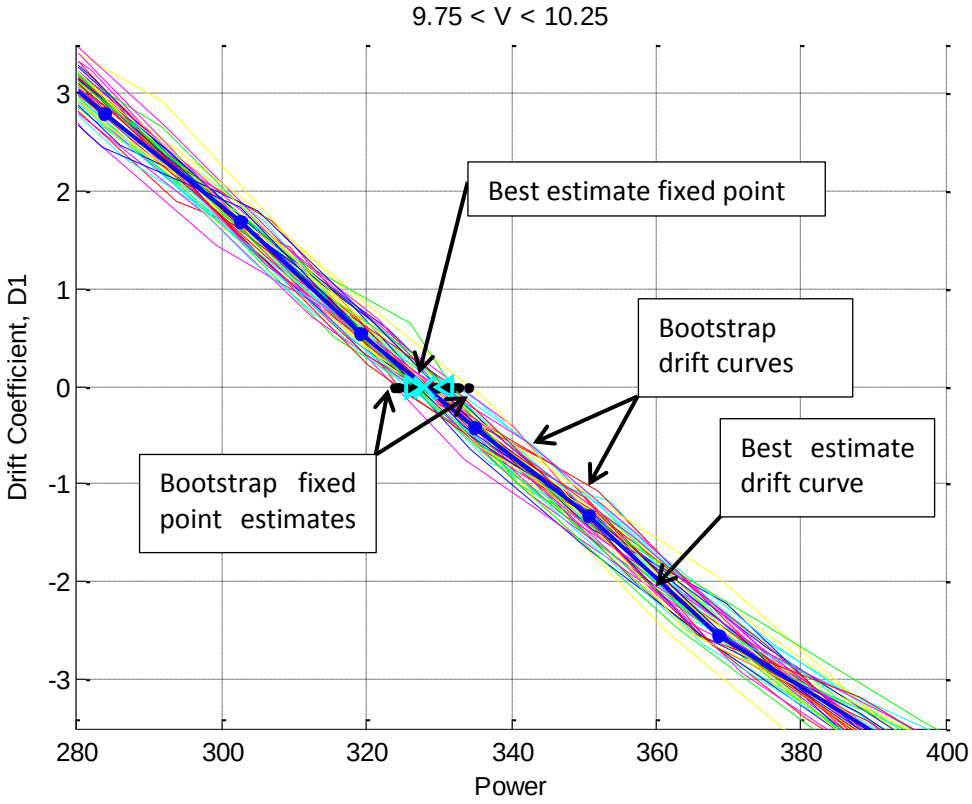

Figure 22 – Drift curve with bootstrap drift curves and fixed point for wind speed bin centered on 10 m/s

4.4 Simulated Langevin Power Curve

To test that the Langevin method had been implemented correctly, a simple drift field reconstruction was performed on some simulated data. The Langevin equation, shown in Equation (15), was used to generate the simulated power data. The drift field was then reconstructed with this simulated data and the fixed points were extracted. Since the Langevin equation was used to generate the data, the Langevin method should be able to give back the inputs used when generating the data.

$$\dot{p}(t) = D_i^{(1)}(p) + \sqrt{D_i^{(2)}(p)} \cdot \Gamma_i(t) \tag{15}$$

A simplified drift field was created to generate the data. The drift and diffusion coefficients were calculated with Equations (16) and (17). As seen in Equation (16), the drift coefficient was simplified to a linearly increasing drift coefficient in each velocity bin proportional to α_i. The diffusion coefficient was simply a constant, β_i, for each velocity bin. These were both simplified further such that all velocity bins had the same α and β. Values of α and β were chosen to be 0.1 and 0.01, respectively.

$$D_i^{(1)}(p) = -\alpha_i\big(p(t) - P_{FP}(u)\big) \tag{16}$$

$$D_i^{(2)}(p) = \beta_i \tag{17}$$

The stochastic component to the equation obeys the following conditions shown in Equations (18) and (19). The autocovariance shown in Equation (19) is "δ-correlated" where δ is the Dirac delta function indicating that each point is independent with a standard deviation of $2\delta_{i,j}$. The random

component was assumed to be normally distributed. This delta coefficient was also simplified such that all $\delta_{i,j}$ had the same value of 0.01.

$$\langle \Gamma_i(t) \rangle = 0 \tag{18}$$

$$\langle \Gamma_i(t) \cdot \Gamma_j(t+\tau) \rangle = 2\delta_{i,j}\delta(\tau) \tag{19}$$

Finally, to finish generating the drift field, a generic power curve was defined as shown in Equations (20) and (21). The rated speed was set to 12 m/s and the rated power was set to 1.

$$P_{FP}(V) = \begin{cases} aV^3 & \text{for } V < V_{rated} \\ P_{rated} & \text{for } V \geq V_{rated} \end{cases} \tag{20}$$

$$a = \frac{P_{rated}}{{V_{rated}}^3} \tag{21}$$

The resulting power curve and drift field are shown in Figure 23. The drift field is a representation of all of the drift curves in one plot. The hot colours indicate a positive drift and the cool colours indicate a negative drift.

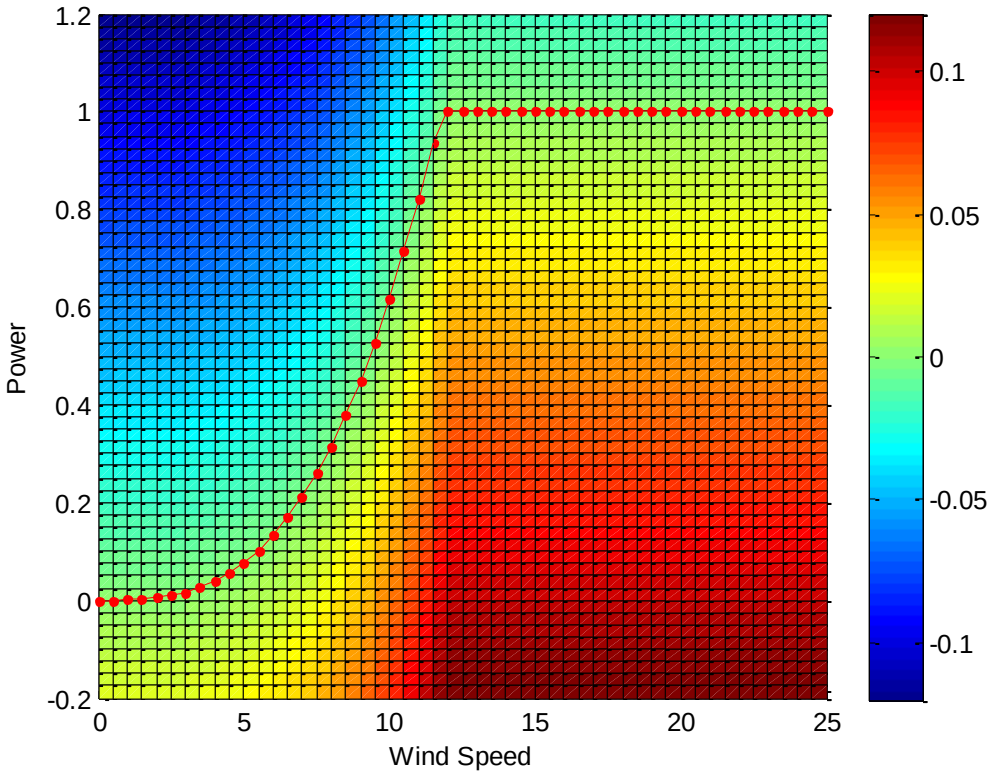

Figure 23 – Drift field of theoretical Langevin wind turbine

To generate the time series, an initial value for the power production was assumed and the Langevin equation was integrated to step through time. The power was set to initially be at the fixed point corresponding to the initial wind speed. Equation (22) was then used to integrate the Langevin equation through time.

$$p(t_{n+1}) = p(t_n) + \Delta t \cdot D^{(1)}(p(t_n)) + \sqrt{\Delta t \cdot D^{(2)}(p(t_n))} \cdot \Gamma(t_n) \tag{22}$$

The wind speed driving the fluctuations of power output was taken from the met mast upstream of the Nordtank turbine at Risø. This wind speed enters into the simulation of power data when calculating drift as shown in Equation (23).

$$D^{(1)}(p(t_n)) = -\alpha\big(p(t_n) - P_{FP}(u(t_n))\big) \tag{23}$$

The resulting simulated power data fluctuates around the fixed point power curve resulting in the point density as shown in Figure 24.

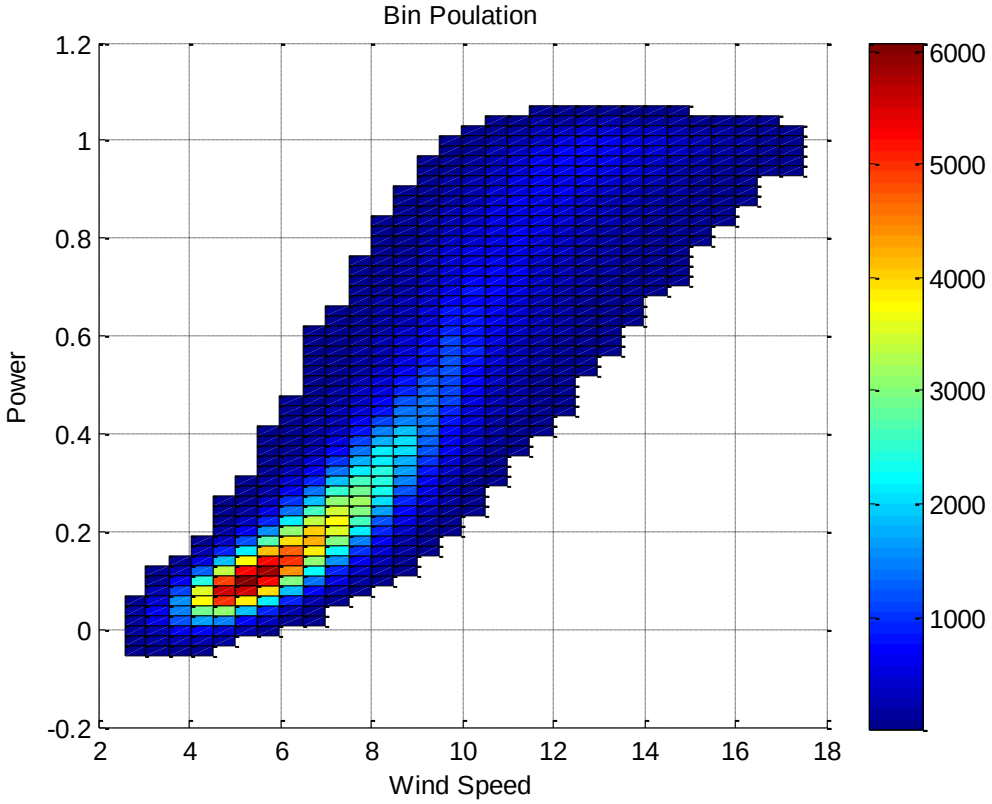

Figure 24 – Simulated power data based on input wind speed

This wind speed and power data was sent to the Langevin method to determine if the fixed point power curve could be estimated accurately from the data.

4.4.1 Langevin Reconstruction of Simulated Power Data

The simulated data generated by the Langevin equation was run through the Langevin method to see if the method could determine the inputs used when simulating the power data. The resulting full drift field is plotted as shown in Figure 25. This figure can be compared with the drift field used to generate the data shown in Figure 23.

When comparing with the actual fixed point power curve used to generate the data, it is noticed that the sharp corner at the rated power is not well reproduced. A comparison between the actual and recreated fixed point power curve is shown in Figure 26. The data was also used to calculate the IEC power curve which was ncluded for comparison. The uncertainty of the power curve is shown by the bounded area around the fixed point power curve estimation.

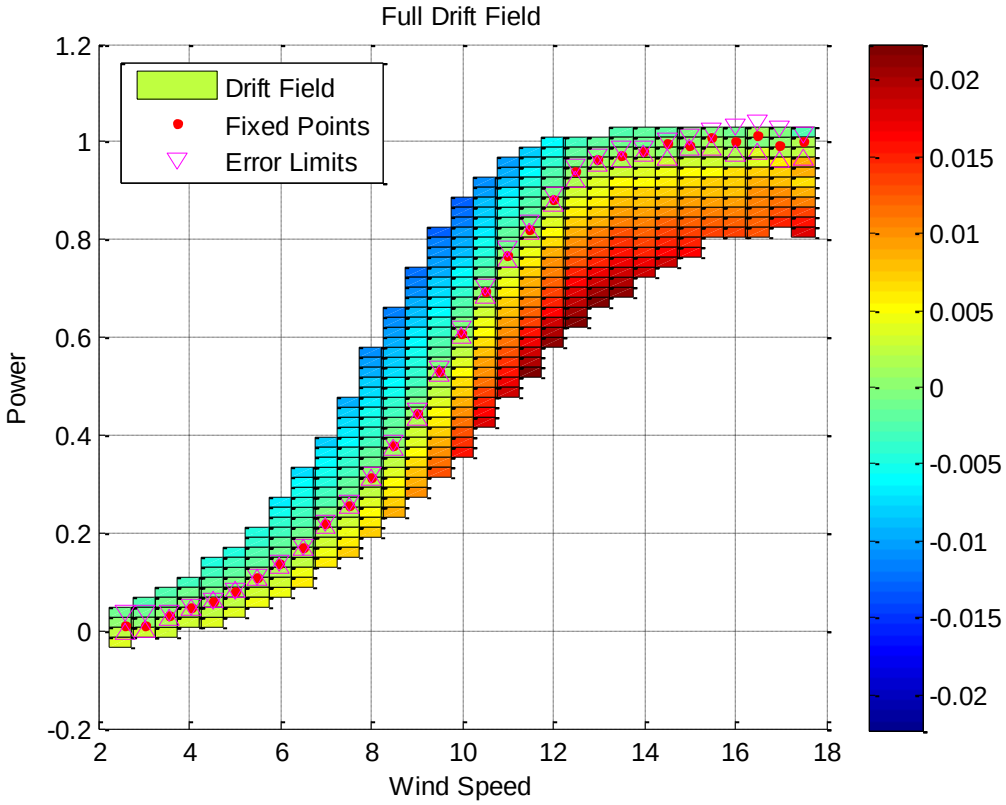

Figure 25 – Drift field from Langevin reconstruction of simulated data

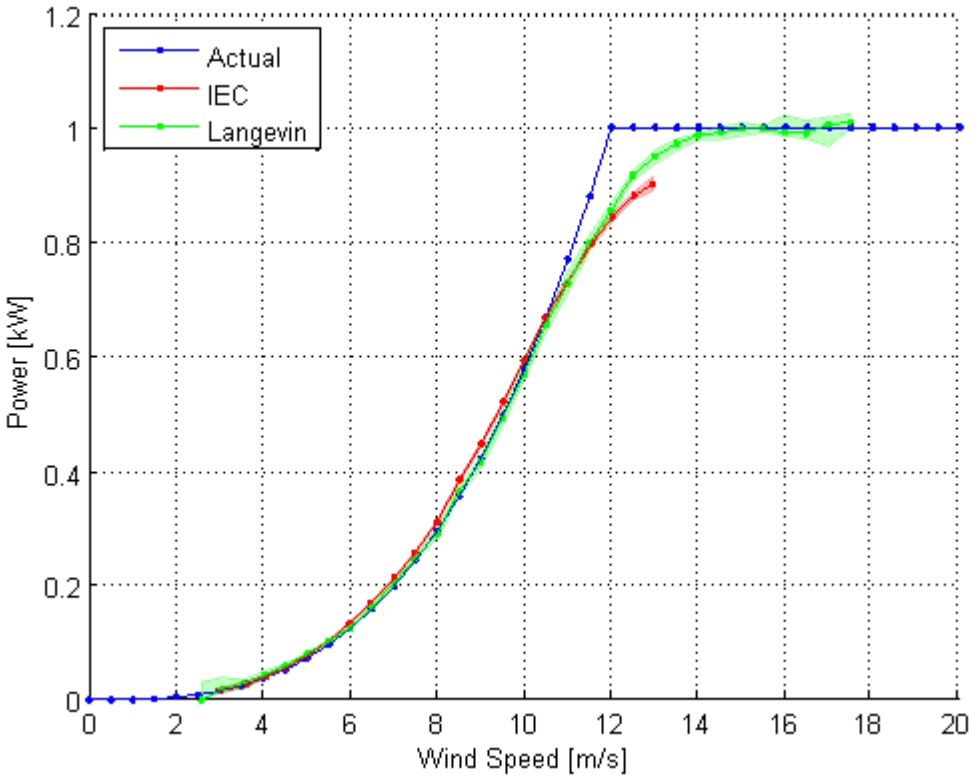

Figure 26 – Comparison of power curves predicted by IEC and Langevin methods

As seen in Figure 26, in most regions of the power curve, the Langevin method accurately reproduces the power curve used to generate the data. The power curve is not accurately reproduced at the region around the transition to the rated power. This is due to fluctuations of the wind speed along with the discrete wind speed bins. This sharp transition is difficult to resolve by the method and this problem has been found in previous work [3]. When the mean wind speed is around the rated power of this fictional power curve, turbulence will cause the wind speed to continuously fluctuate back and forth between the cubic region and the constant region of the power curve confusing the Langevin method. However, a more realistic power curve with a smoother transition would be easier for the method to reproduce. Further investigation into why the knee of the power curve cannot be accurately reproduced using simulated power data is one area suggested for further investigation as described in the recommendation in Section 8.

One of the drift curves was plotted for comparison with the theoretical drift curve used to generate the data as shown in Figure 27. The drift curve predicts a similar fixed point to the actual value; however, the slope of the estimated drift curve is lower than the actual. This is due to the method of fitting procedure for estimating the drift. The drift estimate for the point at a power of 0.24 is shown in Figure 28. The averaged response of the power signal shown in Figure 28 does not exactly match the actual response which would ideally be found from the data. This is because the actual response is calculated assuming the wind speed is not changing with τ. For the averaged response, calculated from many power time series segments as described in Section 4.2.3, the wind speed is changing causing the averaged response to deviate from the actual. A deeper explaination of why the actual response cannot be reproduced by the averaging procedure is recommended for further investigation in Section 8.

Additionally, the actual drift is the slope of the response at $\tau = 0$; however, since we take the slope of the best fit line which is fit over some range of τ, the drift estimate will typically be lower since the response is flattening out with increasing τ.

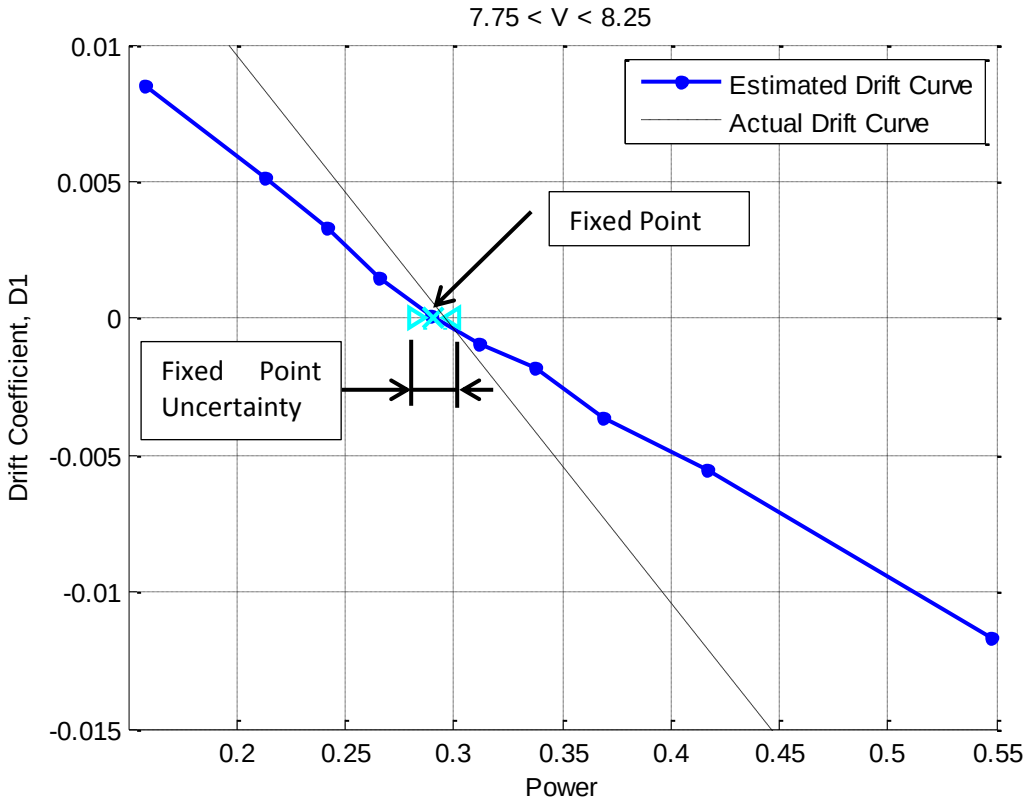

Figure 27 – Comparison of one estimated drift curve with actual

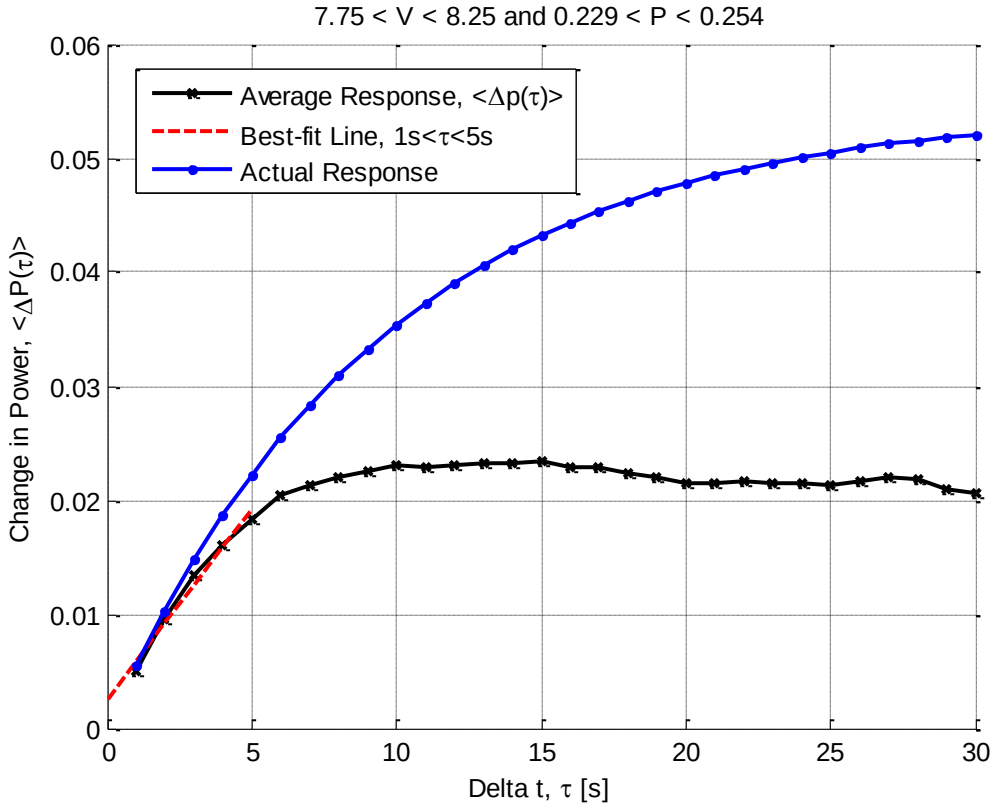

Figure 28 – Comparison of one drift estimate with actual

These results provide some confidence that the Langevin method has been implemented correctly. The problems discussed above have reasonable explanations and Even though the drift curve is not perfectly reproduced, the descriptions above provide some explanation as to why the deviation occurs. These problems are believed to be inherent in the Langevin method and not due to an error in the method's implementation. Since the drift points still have the correct sign and the drift curve still have the same shape, the fixed points are still estimated relatively accurately. The deviations in the fixed points close to the sharp corner of this artificial power curve have also been observed in previous work giving some confidence that it is not the result of an error in implementation.

5 Application of the Langevin Method

In this section, the application of the Langevin method is examined on the measurement data from the Nordtank turbine described in Section 3. The goal of these analyses is to determine if the Langevin method can produce robust, reliable results to determine if it is ready to be applied in industry now or if it needs more development.

Firstly, all of these tests of the Langevin method were performed using an estimate of the driving wind speed calculated with several cup anemometers across the rotor area. However, in past analyses, the shear corrected wind speed has never been used. Additionally, in many test sites, such a well instrumented met mast is not available. Langevin fixed point power curves estimated using the shear corrected wind speed and just the hub height measurement were compared as described in Section 5.2.

The Langevin method has some input parameters which must be chosen by the user which are currently chosen using "engineering judgment". The sensitivity of the Langevin method's outputs to these input parameters was investigated as described in Section 5.3.

Additionally, in previous work [3], it was shown that by pre-processing the wind speed measurements, the outputs of the Langevin method can be improved. The necessity of this pre-processing and the potential reasons why it improves the Langevin results are discussed in Section 5.4.

For this analysis, a reference case was required to compare all of the results to. This reference case is discussed in Section 5.1.

5.1 Base Case

The base case was used as the default case to compare the results to in all of the following cases. The data and inputs for the base case will be described in detail here. In the analyses of the following sections, the ways the Langevin inputs or input data deviate from this base case will be described. All of the other input parameters that are not discussed are the same as this base case.

The measurement data used in the base case is described in Section 3. The data was filtered further such that the 10 minute mean wind direction was between 260° and 310°. This is because the remaining data causes problems for the Langevin method as will be described later in Section 5.4. Wind speed measurements from the full met mast were used to estimate the driving wind speed and help remove some of the effects of wind shear as described in Section 3.3. No further pre-processing was performed on the wind speed measurement. The need for pre-processing of the wind speed will be discussed further in Section 5.4. The power measurements were not corrected for air density since, as discussed in Section 3.4, no reliable temperature and pressure measurements were taken during the period in question.

The Langevin inputs were chosen after testing the sensitivity of each parameter. This sensitivity analysis was presented later in Section 5.3. The measurements were binned based on wind speed using even bin ranges of 0.5 m/s as recommended in the IEC standard. The data was binned again into sub-bins using the power measurements such that each velocity bin had 10 sub-bins with an even number of data points. The interval of time lag, τ, used for fitting to estimate the drift in a sub-bin was [1, 10] seconds, start and end points inclusive.

To estimate the uncertainty of the fixed points, the "Bootstrap Time Series" method described in Section 4.3.3 was used. The bootstrapping process was repeated to obtain 100 estimates of the fixed points. When generating bootstrap time series, data segments of 3000 points were drawn at random and assembled into the new bootstrap time series. In all of the analyses, regardless of the uncertainty estimation method, the uncertainty was expressed as a confidence interval of 90%.

5.2 Shear Correction of Wind Speed

All of the analyses in this report were performed using the shear corrected wind speed calculated as described in Section 3.3. Since many test sites will not have the luxury of several wind speed measurements across the relevant turbine's rotor plane, the effect of using the shear corrected wind speed on the Langevin results was investigated. Additionally, this has not been tried before and in previous work, the Langevin method has been applied to just the hub height wind speed measurement [3].

The data and Langevin inputs from the base case described in Section 5.1 were used with the shear corrected wind speed measurement replaced with only the hub height wind speed. The resulting power curve compared to the base case shear corrected power curve is shown in Figure 29. The range of fixed point uncertainty, using a 90% confidence interval, is shown by the filled area of the same colour as the best estimate line.

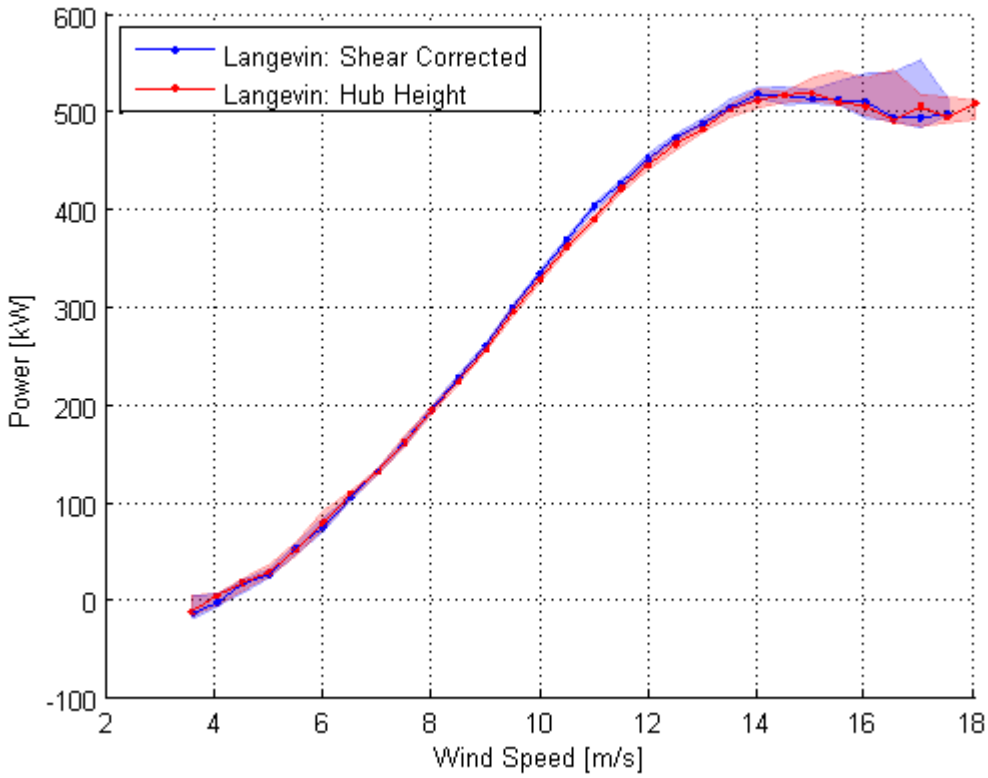

Figure 29 – Comparison of Langevin power curves predicted from shear corrected and hub height wind speed measurements

As seen in Figure 29, the Langevin method is able to predict a similar power curve using either the shear corrected wind speed or only the hub height measurement. This is illustrated further as shown in Figure 30 where the differences between the fixed points of the power curves are plotted. The base case, shear corrected power curve was used as a reference. The envelope of the 90% confidence interval was also shown with a shaded area of the same colour.

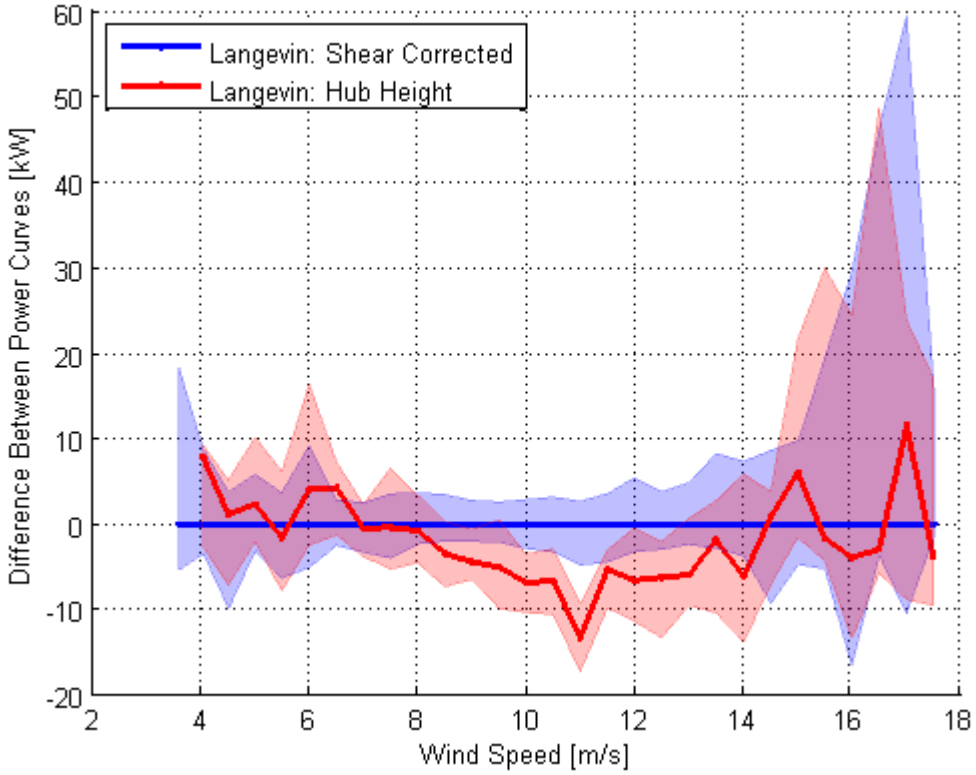

Figure 30 – Difference between Langevin power curves predicted from shear corrected and hub height wind speed measurements

As seen in Figure 30, for a wide range of the power curve, there is very little difference between the two curves. It is possible that the wind shear during this relatively short measurement period was not very severe and was relatively constant. The measurements were all taken during one season (autumn) so seasonal variations would not likely be effectively captured. Some of the fixed points deviate significantly approximately between 10 m/s and 12 m/s. It is possible that there were some periods of strong wind shear at these wind speeds, altering the results slightly.

There are also some secondary benefits of using multiple measurements for estimating the driving wind speed. First, some spatial averaging is achieved in the process. This averages out the very small scale high frequency turbulence which will have likely significantly evolved by the time it reaches the rotor and is too small to affect the turbine and anyways. Additionally, the reliability is improved since we are not relying on a single instrument. If one of the cup anemometers has some error, the effects will be reduced by the other cups.

This shows that a single hub height measurement can be used if that is all that is available. Additionally, if the wind shear is not severe or changing significantly over the measurement period, it can produce reasonably accurate results. However, if the appropriate measurements are available, use of the shear correction is recommended.

5.3 Sensitivity of Langevin Reconstruction to Inputs

Since the Langevin method requires some user defined selections which are currently chosen based on engineering judgment, these choices could potentially be affecting or biasing the results. The effects of these choices on the resulting fixed point power curve were investigated.

5.3.1 Range of τ

Since the range of time lag for fitting during the estimation of the drift points is defined by the user, the sensitivity of the resulting power curve to the user's selection has been investigated. The Langevin method has been used on the data set from the base case described in Section 5.1. The Langevin inputs from the base case were also used, other than the range of time lag for fitting to estimate the drift coefficients. The resulting power curves and uncertainties using τ between [1, 5] seconds, and [1, 10] seconds are shown in Figure 31. The base case uses [1, 10] seconds and was shown in blue.

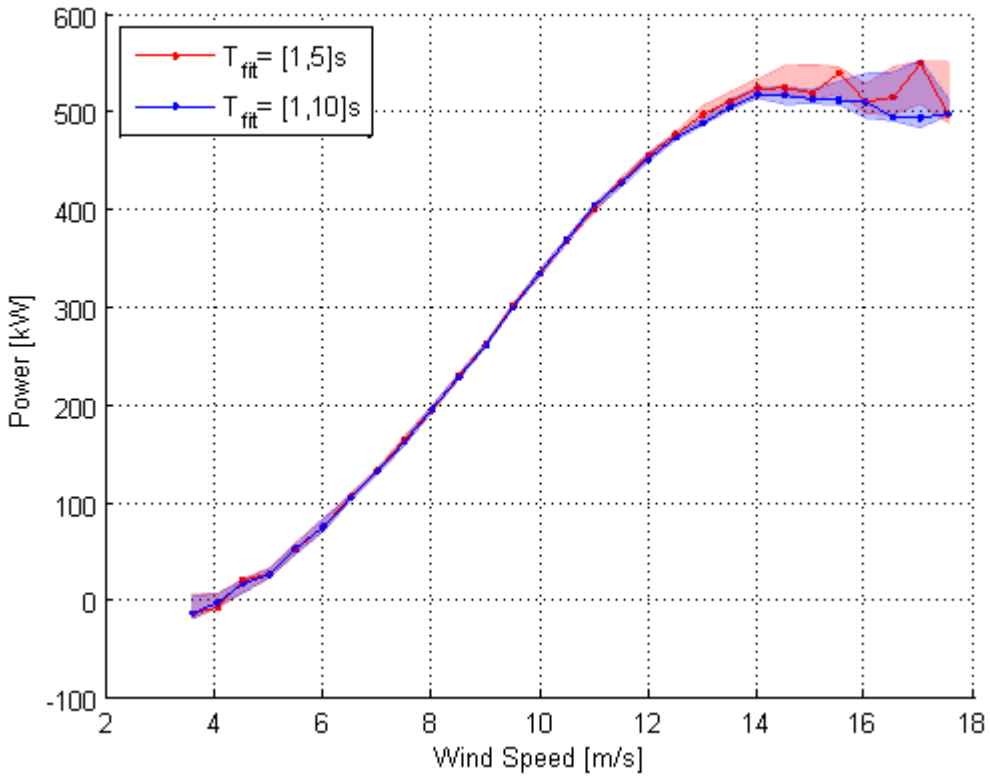

Figure 31 – Power curves predicted using different choices of range of time lag while fitting drift coefficients

As seen in Figure 31, the power curves do not differ significantly. The differences between the fixed points of the power curves are highlighted, as shown in Figure 32, by plotting the difference between the power curves, using the base case power curve with a fitting range of [1, 10] seconds as a reference. The upper and lower bounds of the error bars are also plotted once again using the filled areas.

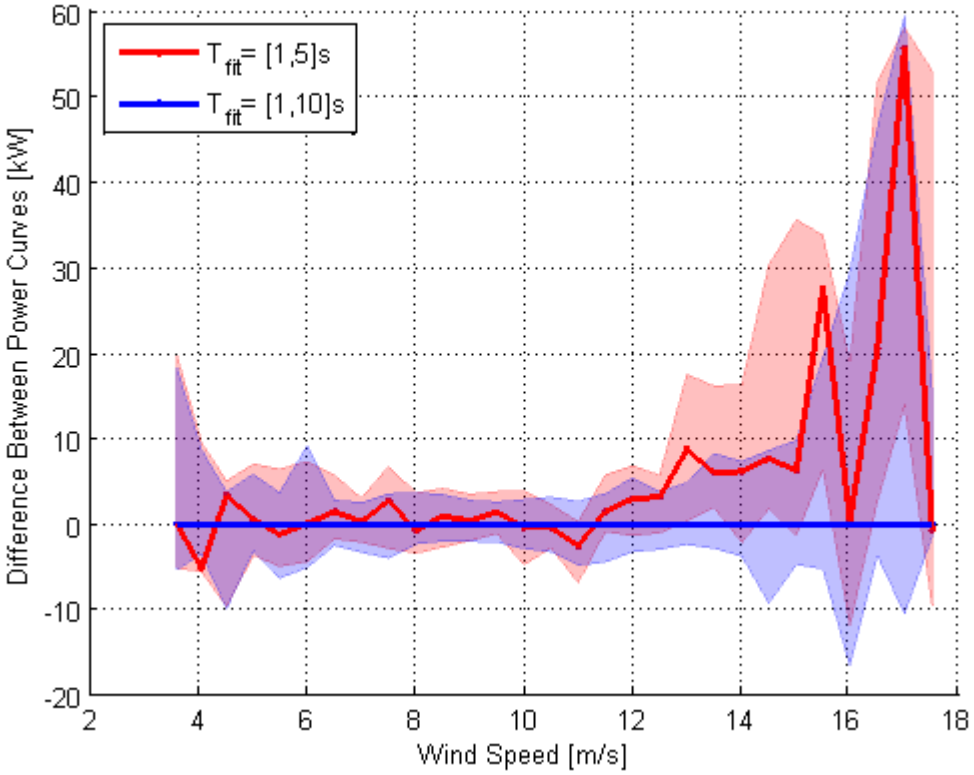

Figure 32 – Difference between power curves with red curve as a reference

As seen in Figure 32, the power curves are close agreement except at the high wind speed bins where there is little data. There is some overlap of the range of the error bars in all of the wind speed bins indicating that the method is somewhat insensitive to the range of time lags chosen within the range up to 10s time lag.

The power curves and difference between the power curves comparing fitting range of [1, 10] seconds with [1, 15] seconds is shown in Figure 33 and Figure 34. The base case is once again shown in blue and was used as the reference power curve for Figure 34.

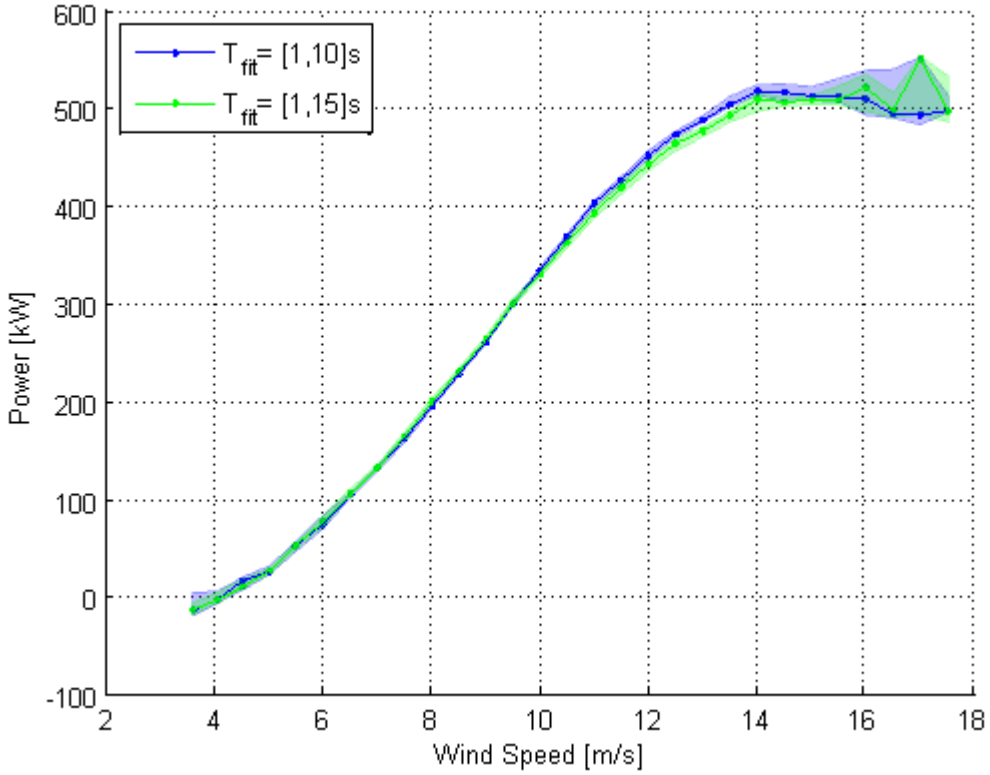

Figure 33 - Power curves predicted using different choices of range of time lag while fitting drift coefficients

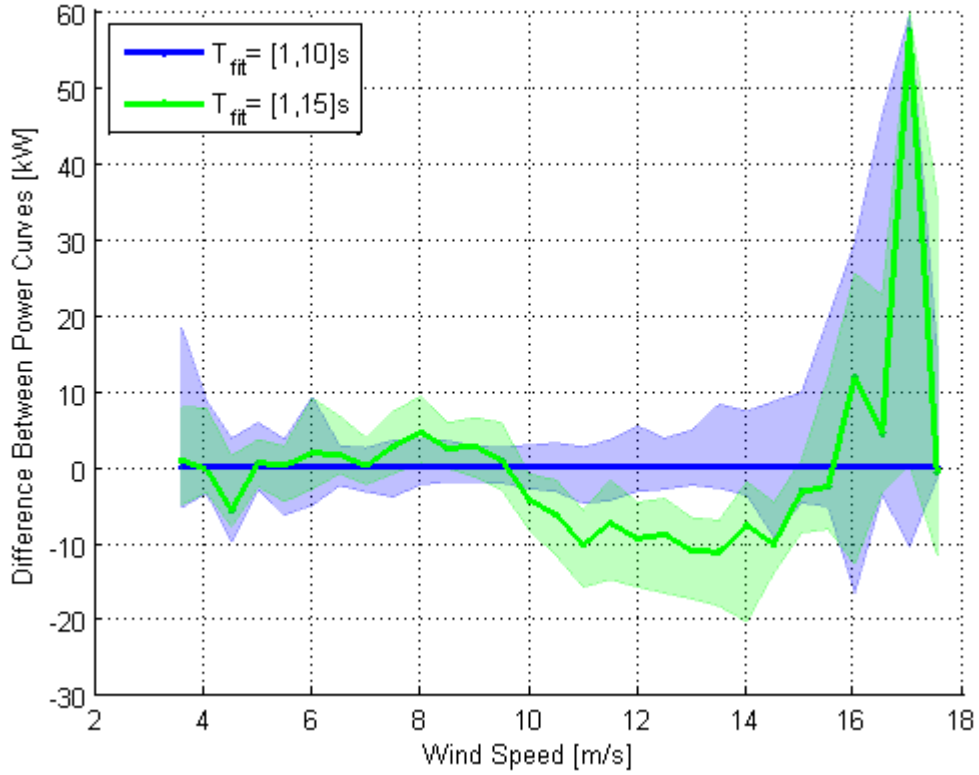

Figure 34 - Difference between power curves with base case as a reference

As show in Figure 33 and Figure 34, the power curves are in agreement at low wind speeds up to approximately 10m/s and then, above this wind speed, the green curve deviates and predicts lower fixed points than the base case. This indicates that the fitting range used can significantly impact the results obtained. The reason for this deviation is demonstrated by looking at the fitting from low to high wind speeds as shown in Figure 35, Figure 36, and Figure 37 plotting the fitting of an example drift coefficient for each of the 6 m/s, 10 m/s, and 14 m/s wind speed bins, respectively. The fitting range from the green curve of [1, 15] seconds was used for fitting to demonstrate why the resulting power curve deviates from the base case.

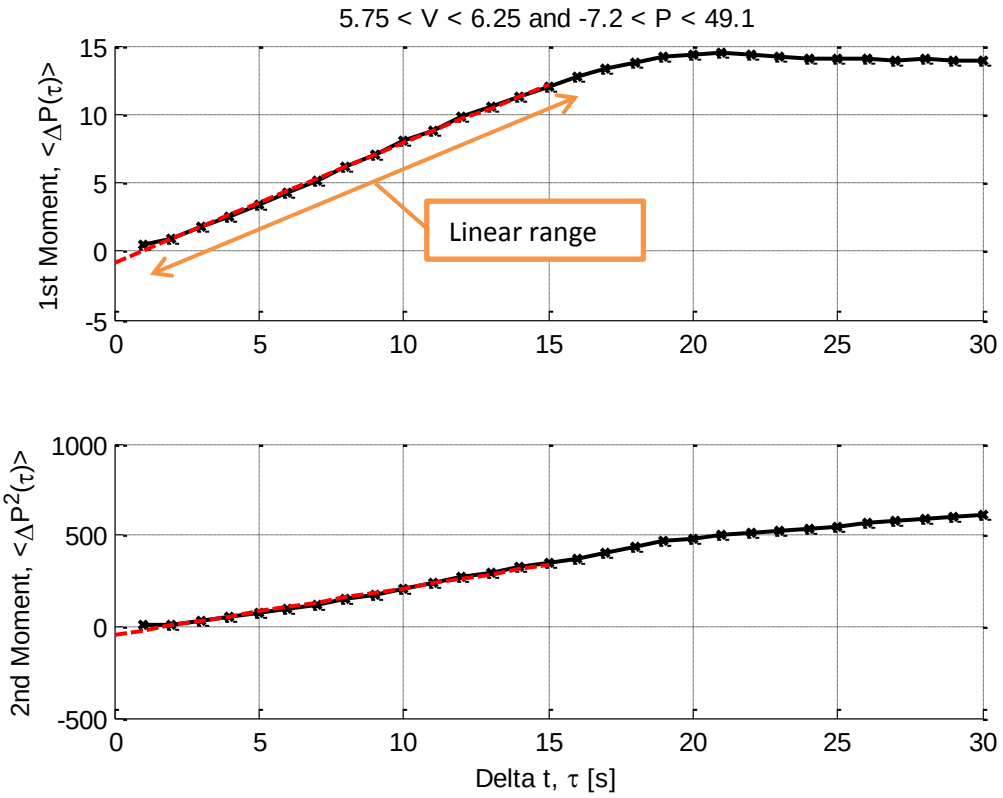

Figure 35 – Fitting of drift coefficient for bin centered on 6m/s

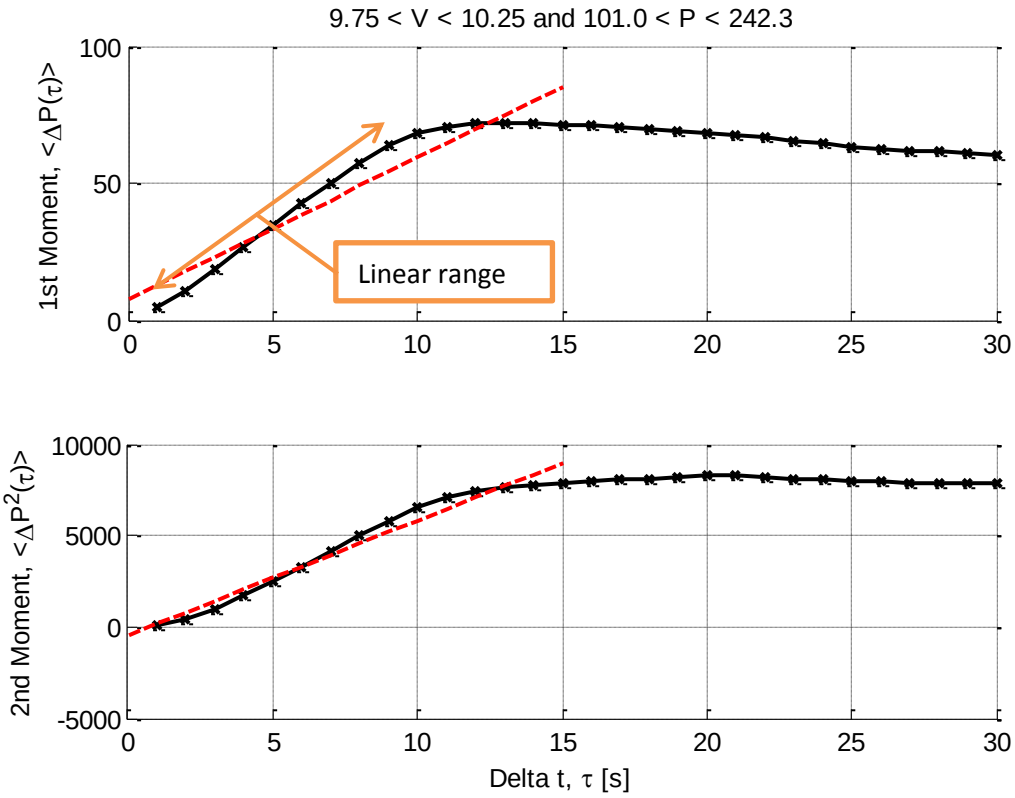

Figure 36 – Fitting of drift coefficient for bin centered on 10m/s

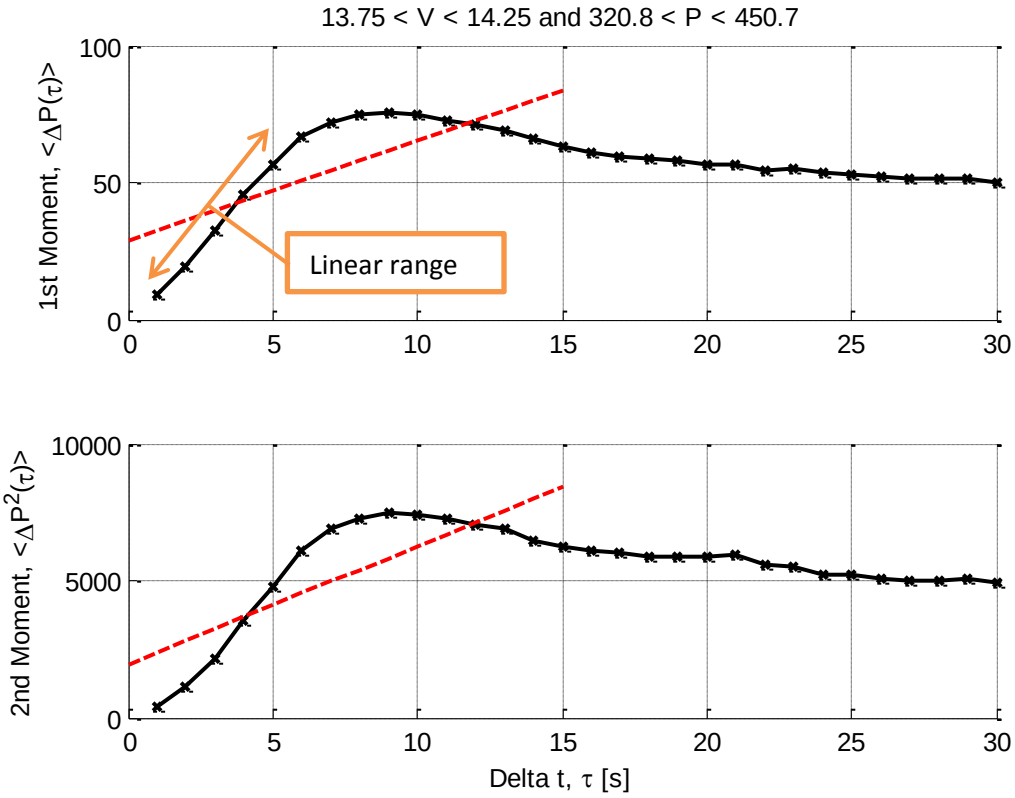

Figure 37 – Fitting of drift coefficient for bin centered on 14m/s

In all three plots above, there is some linear range of the averaged responses where the slope is relatively constant; however, as the wind speed increases, the linear region is shorter and shorter. In Figure 35, any range up to approximately 18s will give a similar approximation of the drift since the dynamic responses has a long linear range. In Figure 37, however, the range of τ used is greatly affecting the predicted drift coefficient. Since the drift is defined as the slope as $\tau \to 0$, and the dynamic response typically flattens out after several seconds, using points outside the linear range will induce errors in the drift curve. The fit shown in Figure 37 uses several points past the end of the linear range; the slope of the fit is underestimating the drift causing the drift curve to flatten out. This is shown in the drift curves from low to high wind speeds show in Figure 38, Figure 39, and Figure 40.

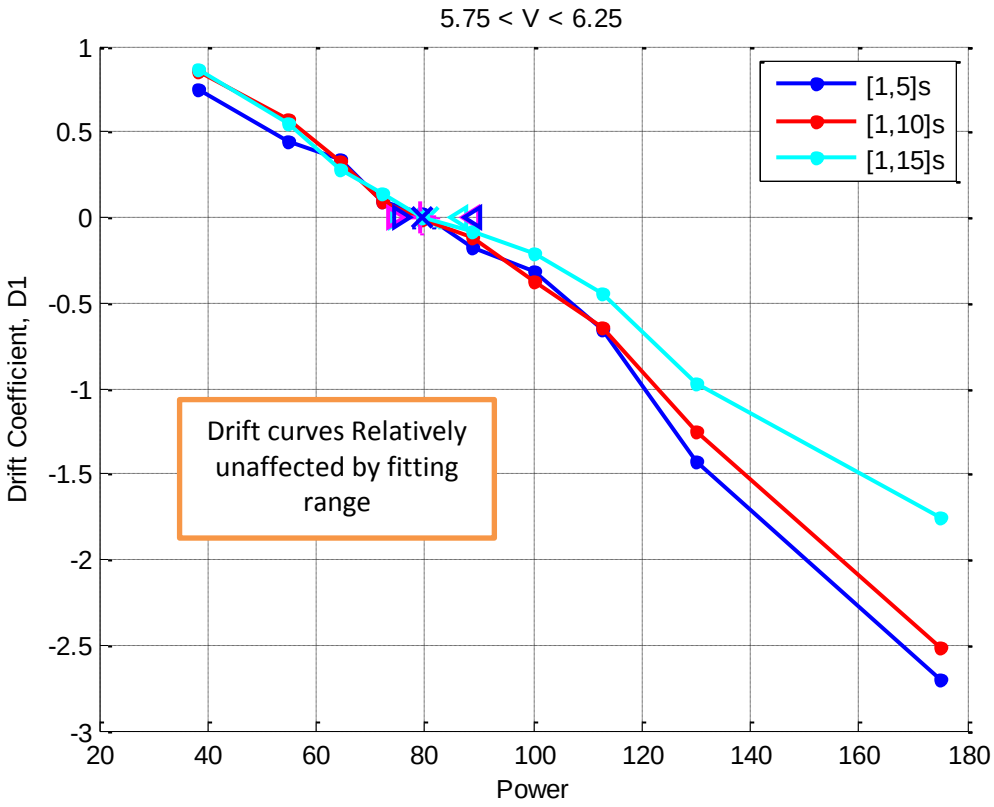

Figure 38 – Drift curves for bin centered on 6m/s

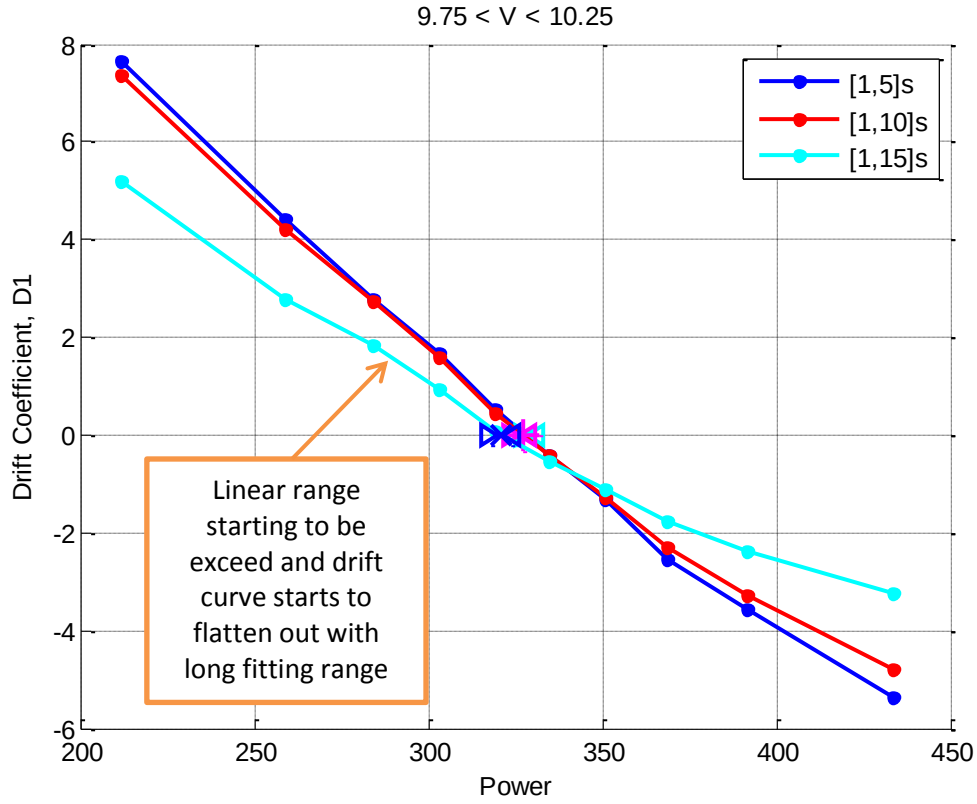

Figure 39 – Drift curves for bin centered on 10m/s

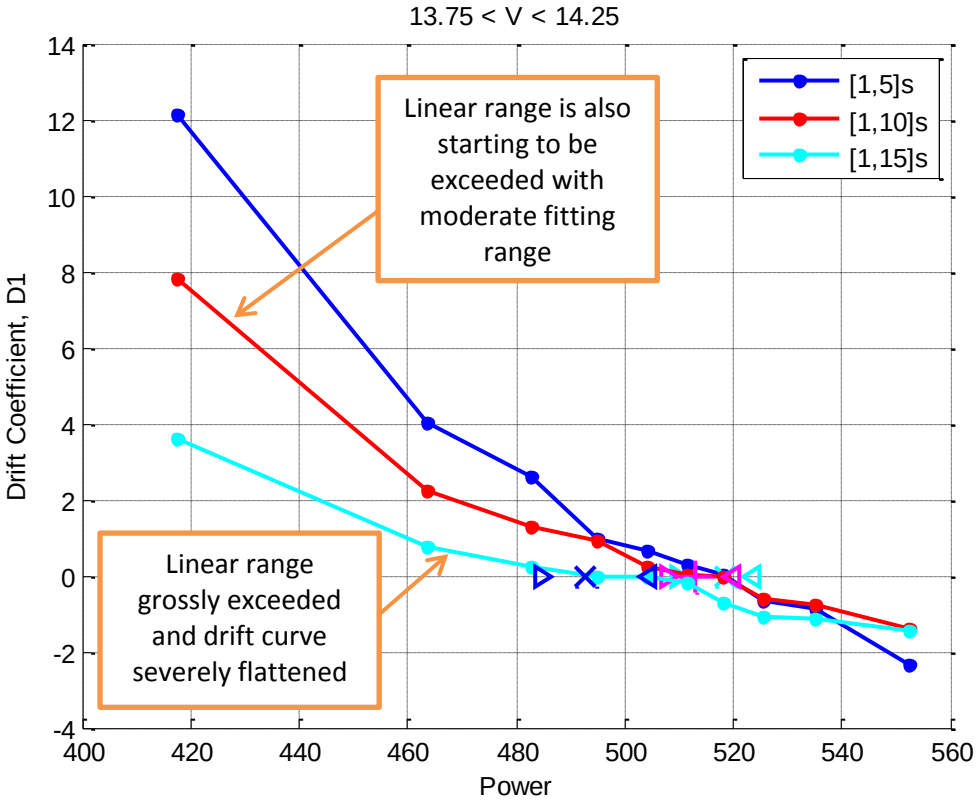

Figure 40 – Drift curves for bin centered on 14m/s

As seen in Figure 38, the drift curves are nearly unaffected by the range of τ used; however, as the wind speed increases, the drift curve is flattened out when using long fitting ranges. This can significantly affect the resulting fixed point as shown in Figure 40.

The resulting drift curves at low wind speeds, and therefore fixed points, are not as sensitive to the fitting range selected by the user since the linear range is relatively large. However, as wind speed increases, the linear range shrinks causing the choice of fitting range to significantly affect the results obtained. Using the current method, all wind speed bins use the same fitting range. With this method, the fitting range should be chosen such that it is short enough not to extend too far past the linear range at high wind speeds as in Figure 37. Additionally, the fitting range should be long enough such that several points of the averaged response are used so noise can be averaged out resulting in a reasonable approximation of the drift.

In future work, the fitting range could be varied for each wind speed bin. This would allow, as much of the linear range to be used as possible to average out errors when estimating the slope of the best fit line. A more sophisticated fitting method could be developed that automatically finds the linear range for each wind speed bin and uses the whole linear range for estimating the drift. The method could vary the fitting range to search for the end of the linear range when the slope begins to significantly change. This reduces the required inputs of the Langevin method and would take some of the guess work out of choosing the fitting range; potentially improving the robustness of the predictions. This suggestion has been included in the recommendations discussed in Section 8.

5.3.2 Power Binning

Since the method of power binning is also chosen by the user, the sensitivity of the results to this selection was analyzed. The data and Langevin inputs, other than the binning method, were taken from the base case described in Section 5.1. The method using even numbers of data points in each bin with 10 and 20 sub-bins was compared with the fixed binning method with a sub-bin size of 10 kW. The resulting power curves and deviation between the power curves are shown in Figure 41 and Figure 42. The base case used the even number of data points method with 10 bins and was used as the reference power curve for Figure 42.

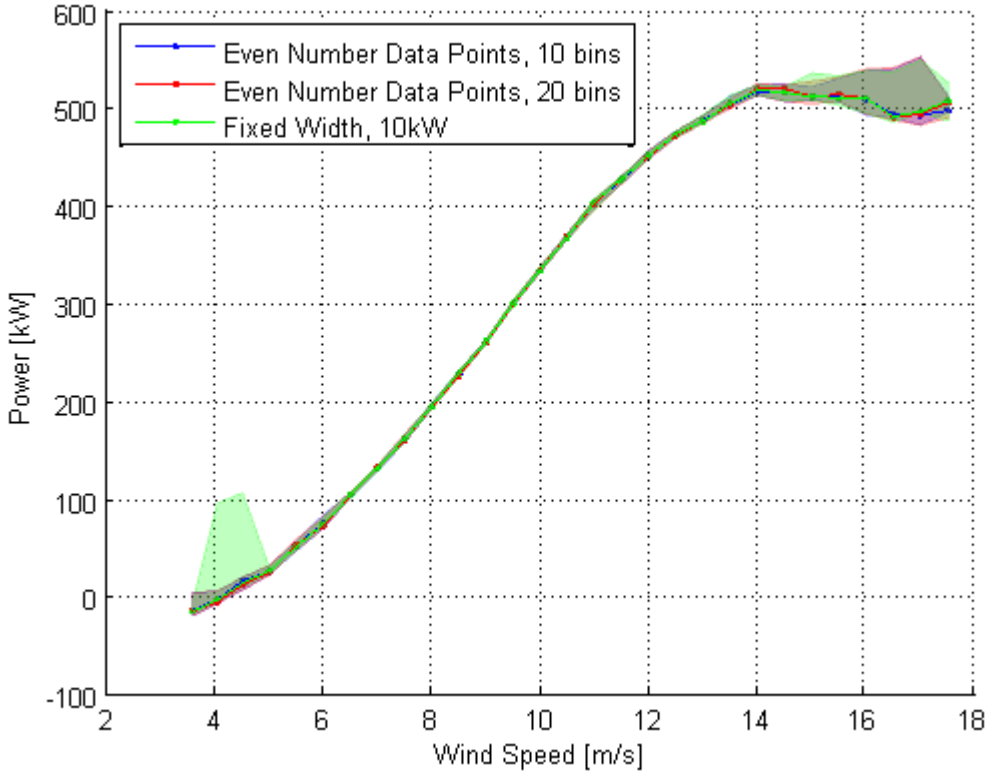

Figure 41 – Fixed point power curves predicted with different methods of power binning

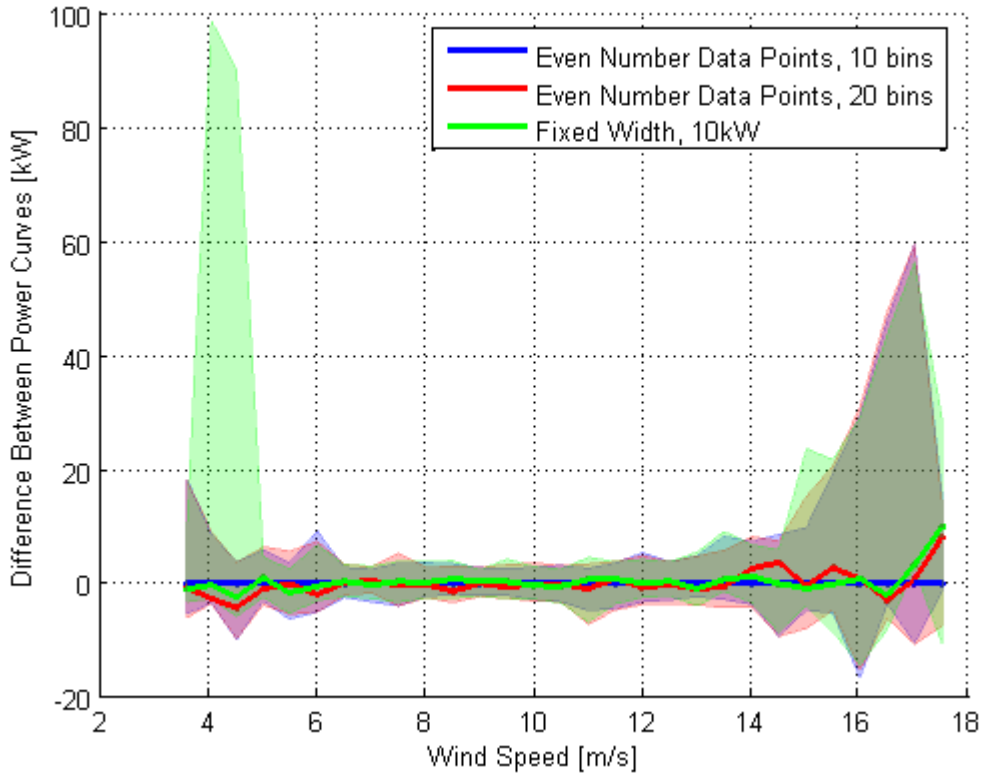

Figure 42 – Difference in power curves predicted with different methods of power binning

As seen in the above figures, all three curves are in very close agreement with nearly identical fixed points. This indicates that the results are somewhat insensitive to the method of binning. This is displayed further by the drift curve for the bin centered on 8 m/s shown in Figure 43. The error bars using the fixed width method are much larger in the wind speed bins up to 5 m/s. The reason is due to the large error of the last sub-bins at either end of the drift curve with small amounts of data in them. This is also illustrated in Figure 43.

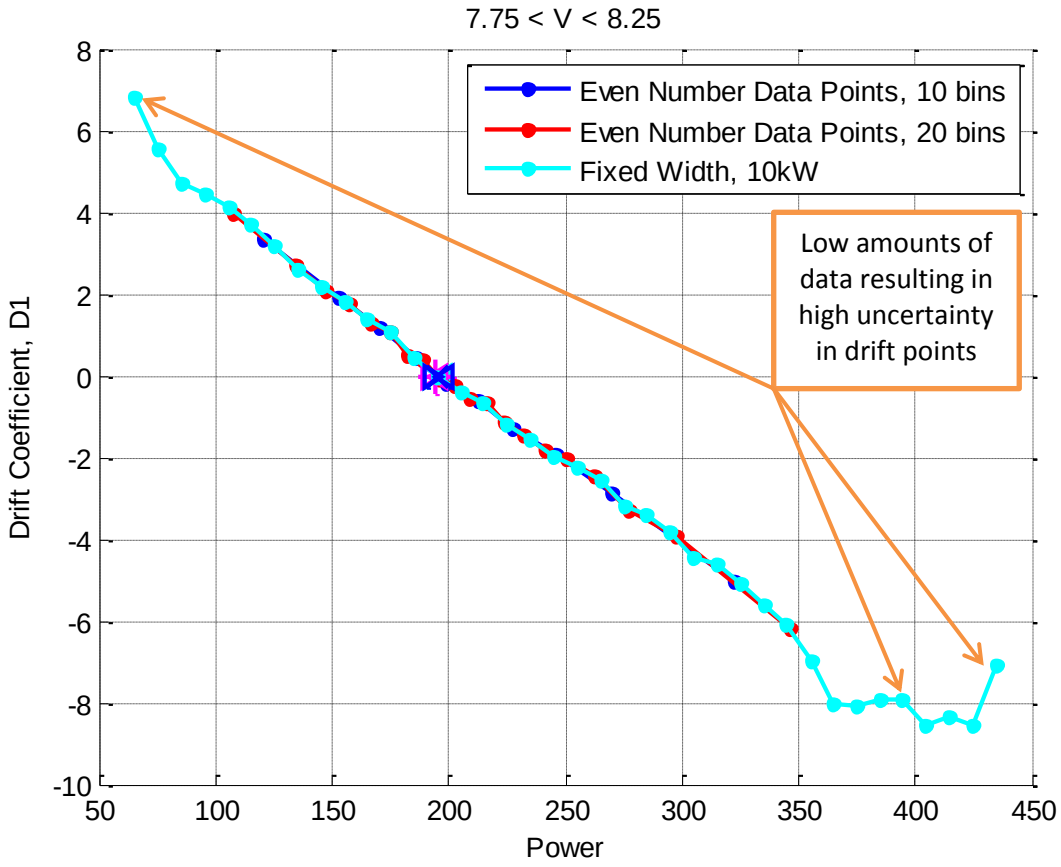

Figure 43 – Drift curve of bin centered on 8m/s bin with different methods of power binning

As seen in Figure 43, the three drift curves and fixed points are nearly identical. However, the fixed width method has additional sub-bins either end of the drift curve. These sub-bins have very little data in them and therefore the uncertainty is very high. This can cause these highly uncertain bins to predict erratic fluctuations in the drift curve at the ends of the curve. In the low wind speed bins of Figure 42, this erratic jumping of the drift curve crossed over the 0 drift line several times causing some of the bootstrap estimate to overestimate the fixed point. To avoid this problem, the method using even numbers of data points in each bin was used as the default method for future analysis.

5.4 Pre-Averaging of Wind Speed

In previous work, some pre-processing of the wind speed data before sending it to the Langevin method has been implemented. Taking a two sided running average over some time period of the wind speed data has been shown to improve the results obtained [3]. This running average is performed as shown in Equation (24).

$$\widetilde{u_T}(t) = \frac{1}{T}\int_{t-T/2}^{t+T/2} u(t)dt \tag{24}$$

This new averaged wind speed is used instead of the raw wind speed when performing the Langevin method. The use of this pre-averaging of the wind speed will be investigated in this section.

In some circumstances, pre-averaging the wind speed does not significantly affect the results as shown in Figure 44. This power curve was generated using the data described in Section 3 with additional filtering on the 10 minute mean wind directions to limit the sector to between 260° to 310°. The Langevin method was applied using the raw wind speed data and pre-averaged wind speed data. The averaging time length, T from Equation (24), was 120 seconds. The Langevin inputs are the same as described in the base case in Section 5.1.

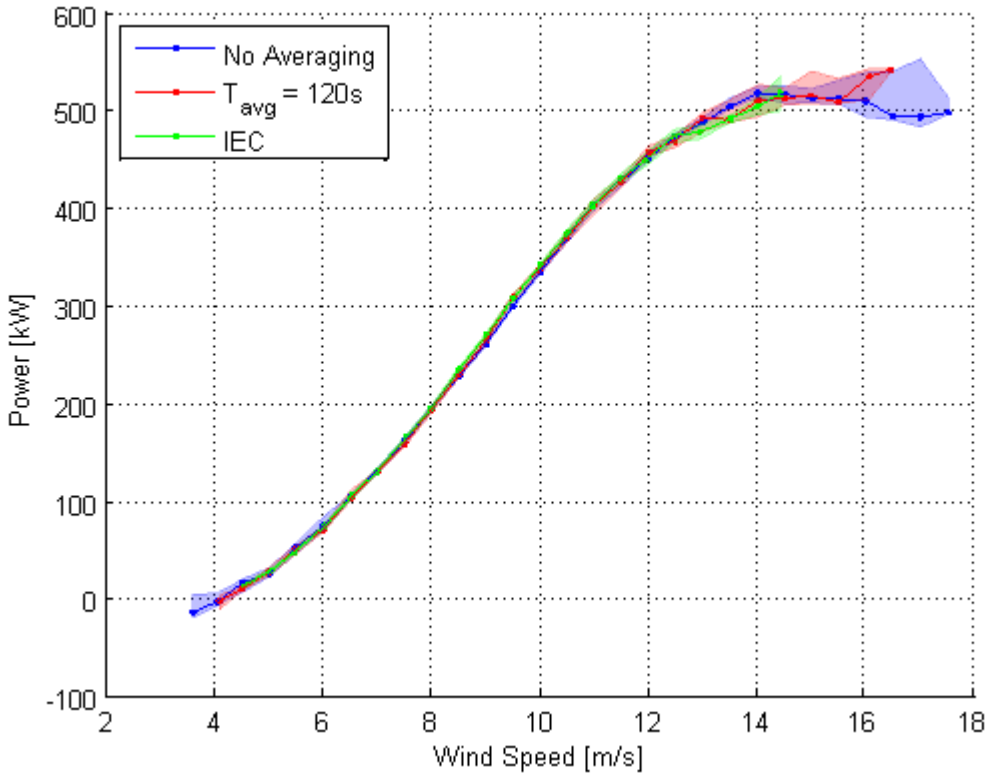

Figure 44 – Comparison of Langevin fixed point power curves with and without pre-averaging for wind directions between 260° and 310°

As seen in Figure 44, the two Langevin power curves do not deviate significantly from each other. Additionally, both power curves are in agreement with the IEC power curve generated using the same data (using raw data without pre-averaging). Using this sub-set of the data, pre-averaging does not have a significant impact on the results.

Using another sub-set of the data, however, pre averaging has a large impact on the Langevin results obtained. The data described in Section 3 was used again with a different sector of between 210° to 260°. The Langevin method was applied using the raw wind speed data and pre-averaged data using the same averaging time length of 120 seconds. The base case Langevin inputs were used once again. The resulting power curves are shown in Figure 45. The IEC power curve was also calculated using this sub-set of the data.

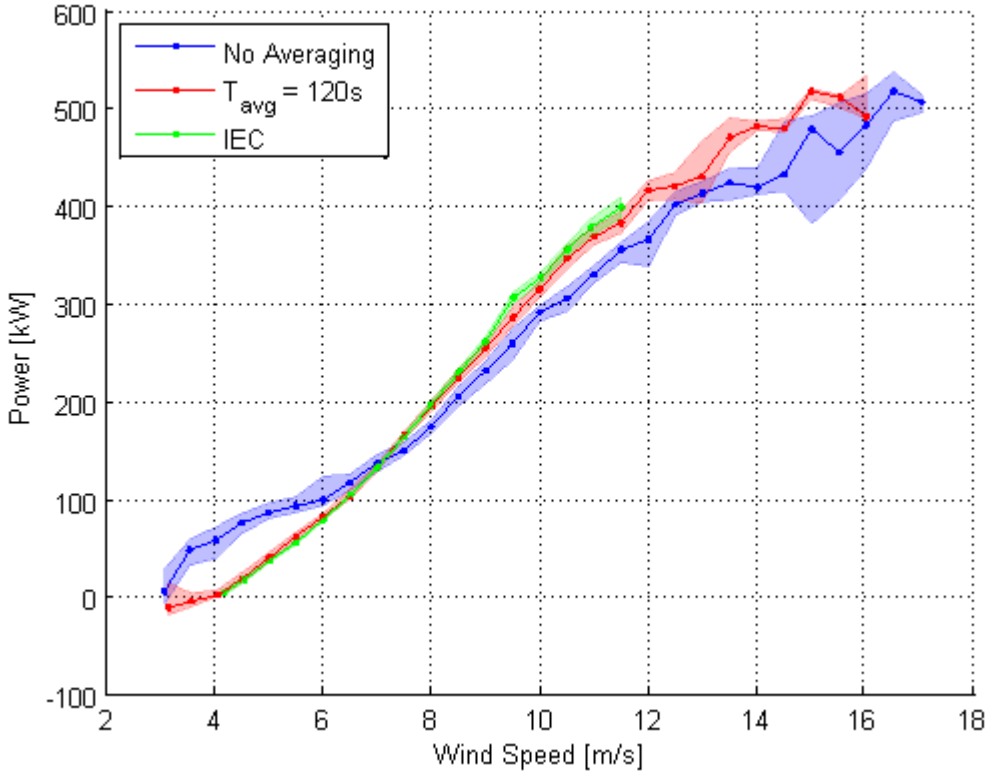

Figure 45 – Comparison of Langevin fixed point power curves with and without pre-averaging for wind directions between 210° and 260°

As seen in Figure 45, the Langevin power curve generated using un-averaged wind speed data is considerably different from both the averaged Langevin and IEC power curves. This un-averaged Langevin fixed point power curve is clearly incorrect as it violates the Betz limit at low wind speeds. Additionally, the un-averaged Langevin results have larger error through the low wind speeds while the averaged Langevin results predict much less error and are in much better agreement with the IEC results. The larger errors at high wind speeds are at least partially due to a lack of data in the high wind speed region, when compared to the power curves shown in Figure 44 from the other sub-set of data.

This large deviation indicates that this Langevin method may not be as robust as the IEC method since this simple change of sector can alter the results by so much. Some pre-averaging of wind speed may be necessary to improve the robustness of the Langevin method. However, this requires the selection of an averaging time length (i.e. T from Equation (24)) so some method of determining this length is required. A method for choosing how to pre-average the data will be discussed later but first, a potential reason for why this deviation occurs is discussed in Section 5.4.1.

5.4.1 Effects of Wind Direction on Correlation of Wind Speed and Power Signals

Comparing the predicted Langevin power curves using un-averaged wind speeds seen in Figure 44 and Figure 45, they are considerably different. However, both methods use the base case Langevin inputs so the only difference in the procedure of generating these power curves was using data from a different sector of wind directions. The problem may root from the difficulty with power curve measurements discussed in Section 2.2.3 related to the spatial separation of the met mast and the turbine. The two sectors are plotted as shown in Figure 46. The distance between the met mast and

the turbine and the turbine diameter are approximately to scale. The two sectors have been named A and B as shown in the figure.

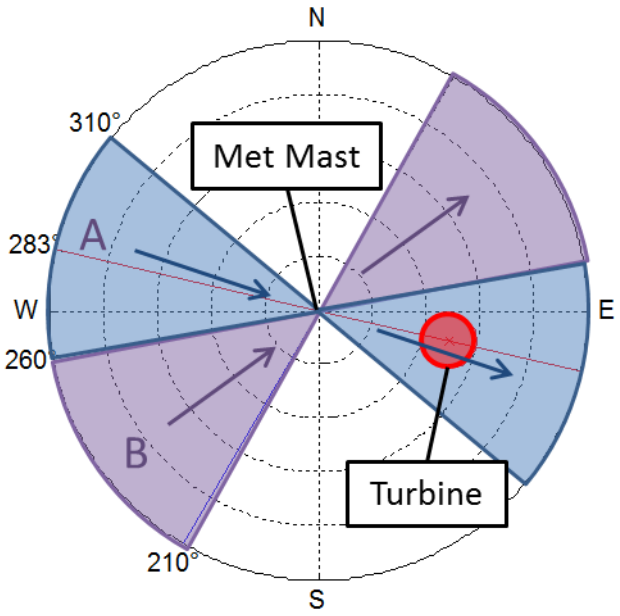

Figure 46 – Visualization of the site layout with the two sectors used for analysis

As seen in Figure 46, the wind coming from sector A will either pass directly over or close to the turbine. The wind coming from sector B, however, will pass some distance from the turbine. Therefore, when comparing the wind seen by the turbine to the wind seen by the met mast, the wind coming from sector A will be more similar than the wind from sector B. This helps improve the correlation between wind speed and power in sector A making it easier for the Langevin method.

Turbulence in the atmosphere contains eddies of many different sizes. Large eddies coming from sector B may be big enough to pass over both the met mast and the turbine, while smaller eddies may only pass over the met mast. Within Sector A more of the eddies measured by the met mast will pass through the turbine's rotor. Additionally, the size of the eddies are related to their frequency. Smaller eddies swirl faster while larger ones swirl slower with a low frequency. With this in mind, the correlation of wind speed and power at many different frequencies can be compared by calculating the coherence of the two signals. An estimate of the magnitude squared coherence of wind speed and power for sectors A and B is shown in Figure 47.

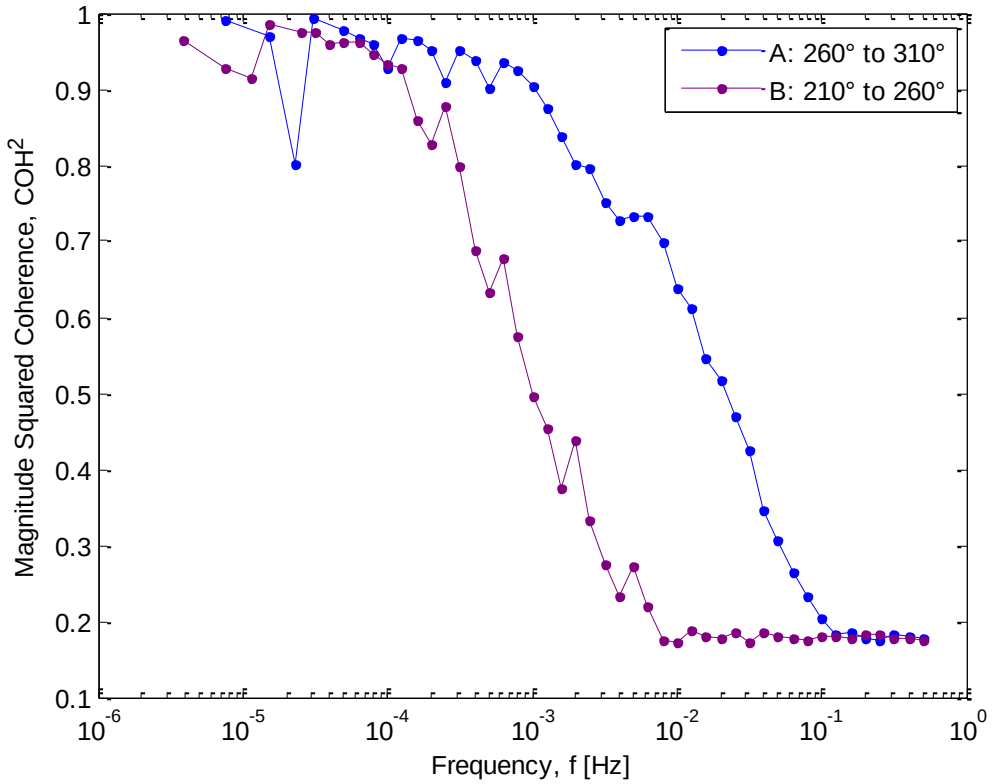

Figure 47 – Comparison of magnitude squared coherence of wind speed and power between Sectors A and B

As seen in Figure 47, as the frequency increases, the correlation between wind speed and power decreases. The larger low frequency eddies have more inertia so they will evolve slower meaning they will not change as much from when they are measured at the met mast to when they pass over the turbine. Additionally, since they are large, they will affect a larger portion of the rotor area. These effects help keep the correlation of these large, low frequency eddies high. The small eddies, on the other hand, will evolve quickly and will not cover a large portion of the rotor, or may miss the turbine completely causing the correlation to drop as frequency increases.

Also, as can be seen in Figure 47, the correlation within sector B drops much sooner at a lower frequency than the correlation within sector A. This could be attributed to more of the eddies missing the turbine as discussed above. For the eddies to induce some change in power on the turbine which can correlated to the measured wind speed, they will have to be large enough to cover both the met mast and some portion of the rotor area. Therefore, although the turbulence characteristics at both the wind turbine and the met mast may be the same, a much larger portion of the turbulence measured at the met mast will not correlate with fluctuations in power in sector B. The effects of this decreased correlation on the Langevin results will be discussed in Section 5.4.2.

5.4.2 Effects of Correlation of Wind Speed and Power on Langevin Method

When investigating the coherence plot for sector A shown in Figure 47, there is some range of frequencies with low correlation between wind speed and power. This low correlation does not seem to impact the Langevin results since, as shown in the power curve in Figure 44, the Langevin method produces sensible results. Investigating the spectrum of the wind speed, shown in Figure 48, the power spectral density, and therefore amplitude of wind speed fluctuations, is decreasing with increasing frequency. Therefore the amplitude is relatively low for the uncorrelated high frequency components of the wind speed fluctuations. The only thing the wind speed measurement is used for

in the Langevin method is to determine which wind speed bin the measurement belongs to. These uncorrelated, high frequency wind speed fluctuations are too small to significantly impact which wind speed bin a particular data point will fall into so the Langevin method can still function.

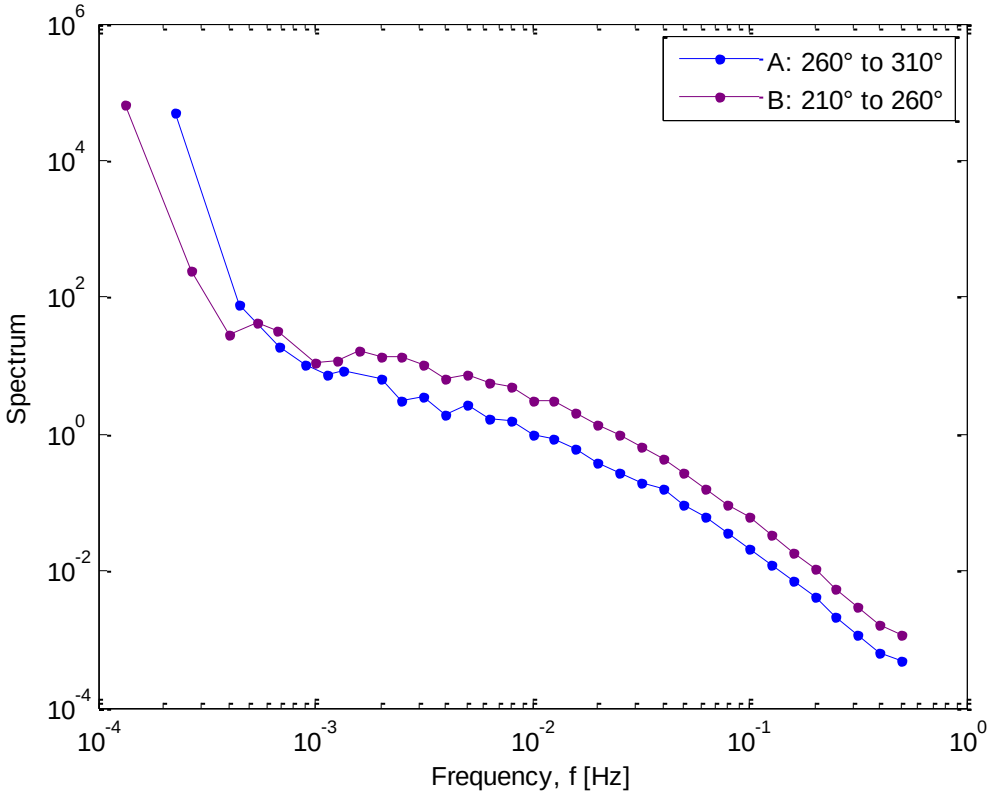

Figure 48 – Comparison of spectra of wind speed from sectors A and B

As seen in Figure 48, the power spectral density of the wind speed from sector B is significantly higher across a wide range of frequencies. This might be attributed to differences in the upstream terrain in the two sectors. Wind in sector A is coming into shore from the Roskilde Fjord, while wind in sector B is over land with more obstacles and higher roughness.

In addition to this increase in turbulence level over all frequencies in sector B, a larger portion of the wind speed fluctuations are uncorrelated with power fluctuations as shown in Figure 47. This combination of larger fluctuations and larger range of uncorrelated fluctuations means that the amplitude of the uncorrelated fluctuations are large enough to effect which wind speed bin a data point gets placed in during the Langevin method. The result is that a given wind speed bin will include data which truly belongs in that bin along with data from the adjacent bins. This can greatly throw off the Langevin method and increase the error of the drift estimates and fixed points as is shown in Figure 45. A possible explanation of why the running mean of the wind speed helps solve this problem is included in Section 5.4.3.

5.4.3 Effects of Taking Running Average of Wind Speed Data

The running average of the wind speed acts like a low pass filter removing all frequency components above some cut-off frequency. This can be seen by comparing the spectra of wind speed before and after taking the running mean as shown in Figure 49 (a). The figure shows the spectra from sector B before and after the running mean has been taken with an averaging length of 120 seconds. Additionally, the gain of the filter function was estimated by taking the ratio between the two spectra as shown in Figure 49 (b).

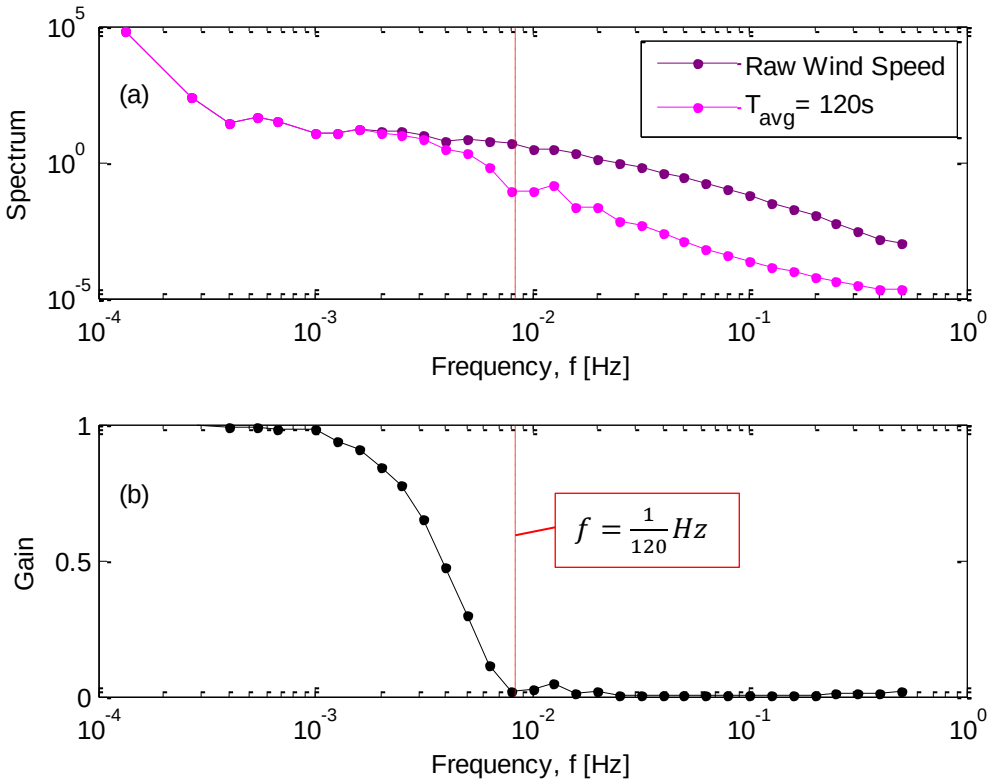

Figure 49 – Comparison of spectra of wind speed with and without running mean of length 120 s on data from sector B

As seen in Figure 49 (b), taking the running mean of wind speed removes all frequency components with periods shorter than the averaging length. There is also a transition region where the gain drops from around 1 to nearly 0. The effect of this filtering of the low frequency wind speed components on the coherence of wind speed and power is shown in Figure 50.

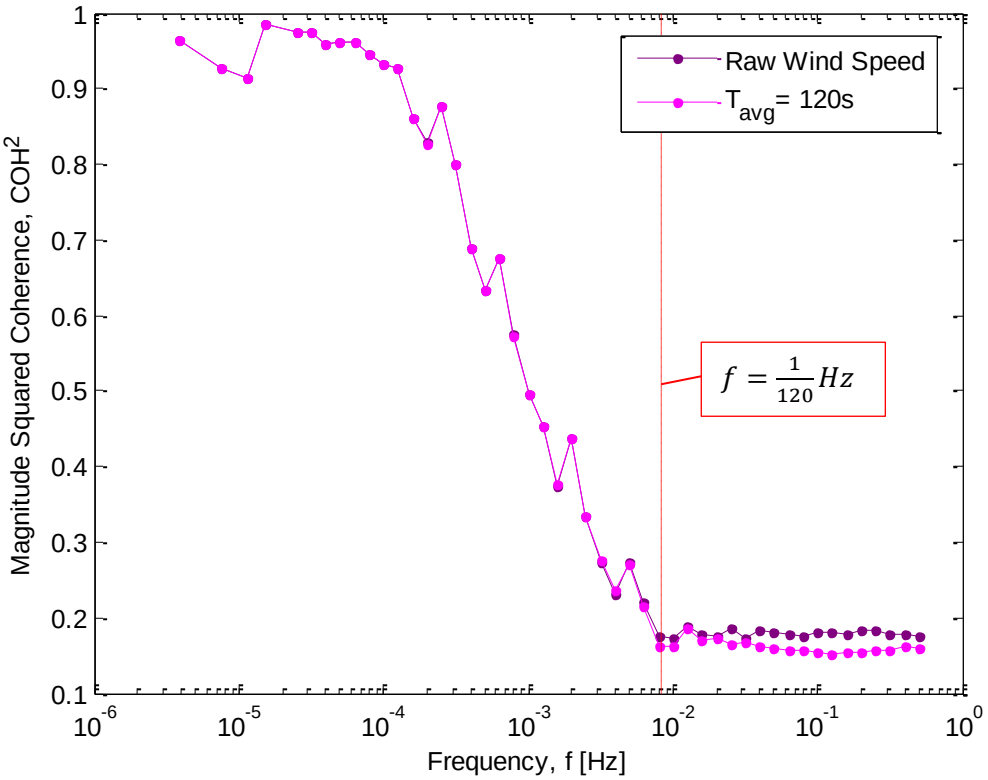

Figure 50 - Comparison of coherence of wind speed and power with and without running mean of length 120 s on data from sector B

As seen in Figure 50, since the high frequency wind speed fluctuations are already uncorrelated to power fluctuations, removing them does not significantly affect the coherence in this region. In the transition region of the filter function, the coherence is altered slightly since the amplitude of the wind speed fluctuations are reduced while the power fluctuations have not been changed. This shows that the uncorrelated high frequency wind speed fluctuations can be removed while maintaining the correlation of wind speed and power at low frequencies.

These uncorrelated, high frequency wind speed fluctuations caused by turbulence measured by the met mast that do not affect the turbine can be thought of as noise on top of the underlying wind speed signal. Furthermore, the uncorrelated high frequency power fluctuations caused by turbulence which flows over the turbine but is not measured by the met mast can be thought of as noise on top of the underlying power signal. The Langevin method can average out this noise on top of the power signal since it performs averaging on the time responses for each sub-bin as described in Section 4.2.2. However, the basic Langevin method does not deal well with this noise on the wind speed. This running average method provides some way of removing the uncorrelated noise on the wind speed signal before sending it to the Langevin algorithm.

This running average can also be applied to sector A where the wind speed and power are better correlated. Figure 51 compares the coherence of wind speed and power of the raw and running averaged wind speed with an averaging length of 120 seconds.

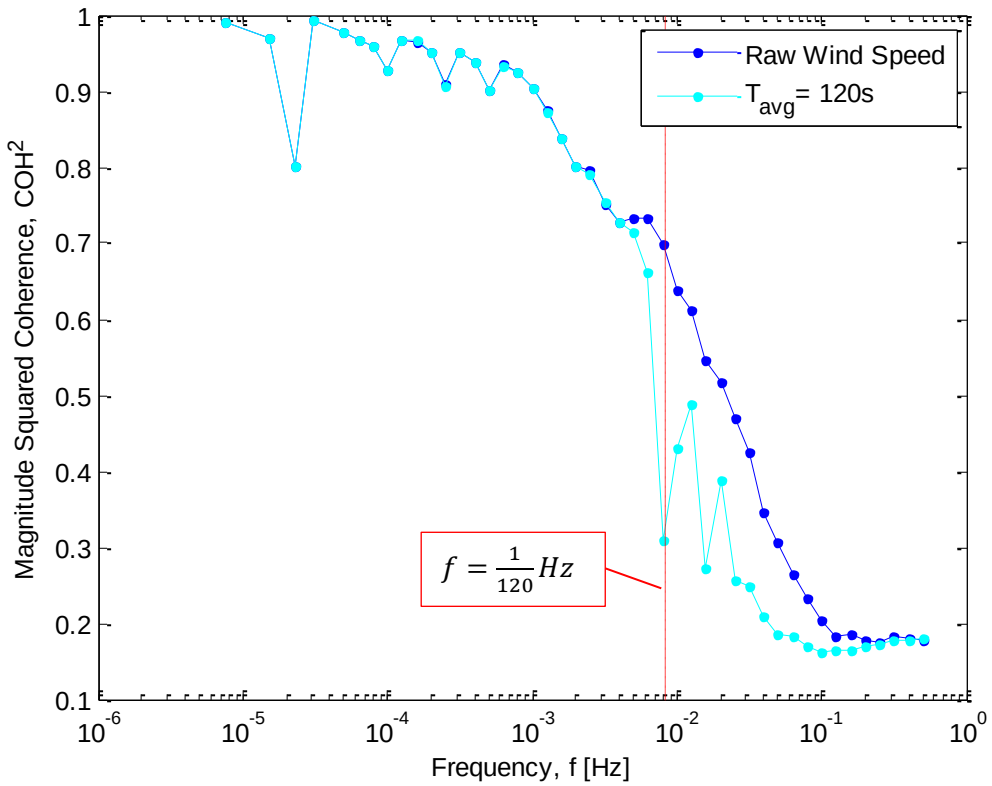

Figure 51 – Comparison of coherence of wind speed and power with and without running mean of length 120 s on data from sector A

As seen in Figure 51, the correlation of the low frequency components of wind speed and power are maintained. However, since an averaging length of 120 s was used, frequency components with periods of less than 120 s which were moderately well correlated were negatively affected. Even though this averaging has damaged some of the correlation between wind speed and power, the Langevin method still predicts a similar power curve as shown in Figure 44. This is likely because the data is originally very well correlated so even though some of the correlation is damaged, it is still able to determine a sensible estimate of the fixed point power curve. That being said, it is not recommended that well correlated portions of the coherence plot are averaged as this risks introducing some uncertainty.

A comparison of the Langevin fixed point power curves predicted from sectors A and B after averaging the wind speed with a 120 second averaging is shown in Figure 52. Additionally, the IEC power curve was made for comparison using data from both sectors.

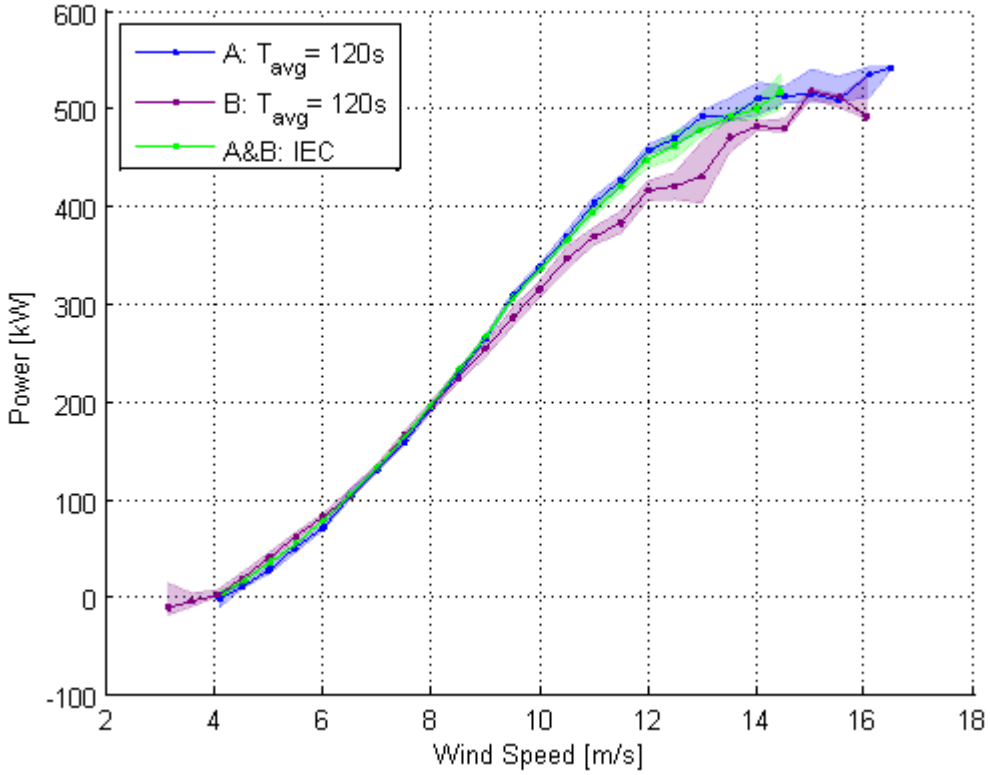

Figure 52 – Comparison of fixed point power curves with standard power curve

As shown in Figure 52, the Langevin method produces more sensible results using data from sector B which agree much better with the IEC power curve and the Langevin power curve from sector A. The power curve is not identical; but, since the correlation between wind speed and power in this sector is much worse to begin with, it cannot be expected to match perfectly. Additionally, the high wind speed bins are much more sparsely populated in sector B than in sector A, which likely accounts for some of the deviation and large error at high wind speeds. Without this running average of the wind speed data, the results shown in Figure 45 are obtained and the data cannot contribute to the Langevin power curve since it confuses the Langevin method. A method for combining data from different sectors will be discussed in Section 5.4.4.

5.4.4 Incorporating Wind Speed Filtering into Langevin Method

As explained in the previous sections, performing some pre-processing of the wind speed data can help improve the Langevin results obtained in poorly correlated sectors. This pre-processing could be included as a permanent step to the Langevin method to improve the results obtained. A method for incorporating the pre-processing of wind speed data is described in this section.

Since taking the running average of the wind speed is essentially just acting as a low pass filter, this could be replaced with a low pass filter that has been designed specifically for the data. The running average's gain, shown in Figure 49, has only one parameter to control its shape; the averaging time period. By using a filter instead of the running average, the shape of the filter function's gain could be much more precisely controlled. A high order filter could be used to obtain a sharp drop in gain above the cut-off frequency. The cut-off frequency could be set based in the coherence plot such that the uncorrelated wind speed fluctuations could be removed while preserving as much of the good correlation between wind speed and power as possible.

As discussed in Section 5.4.1, the wind direction can affect the coherence of wind speed and power fluctuations at different frequencies. Additionally, by filtering out high frequency fluctuations from

all sectors equally, although it could remove noise from the poorly correlated sectors, could damage some of the potentially useful correlation in the well correlated sectors. For this reason, a filter could be designed for each wind direction rather than a single filter for all the data. The data could be split into wind speed sectors, and a filter could be designed for each one. The cut-off frequency of the filter could be chosen based on the coherence of wind speed and power for that sector. After filtering the data in each sector, the data could be reassembled and sent to the Langevin method.

A simple example of this procedure was performed using the same data from sector A and B again. Data from the two sectors was filtered separately using a filter appropriate for each sector. The data was then recombined and sent to the Langevin algorithm.

For simplicity, the running average wind speed was used again to filter the data instead of designing a high order low pass filter as discussed above. When looking at the coherence of wind speed and power of the two sectors, as shown in Figure 47, the coherence drops off at different frequencies. The averaging length was chosen to remove the frequencies with coherence below approximately 0.2, where the coherence plots appear to stop decreasing and plateau. An averaging length of 15 seconds was chosen for sector A while a length of 120 seconds was chosen for sector B. After performing the running averages, the data was re-combined into one time series and was put through the Langevin method. The resulting power curve is shown in Figure 53 with comparison to the IEC power curve generated using the same raw data from sectors A and B.

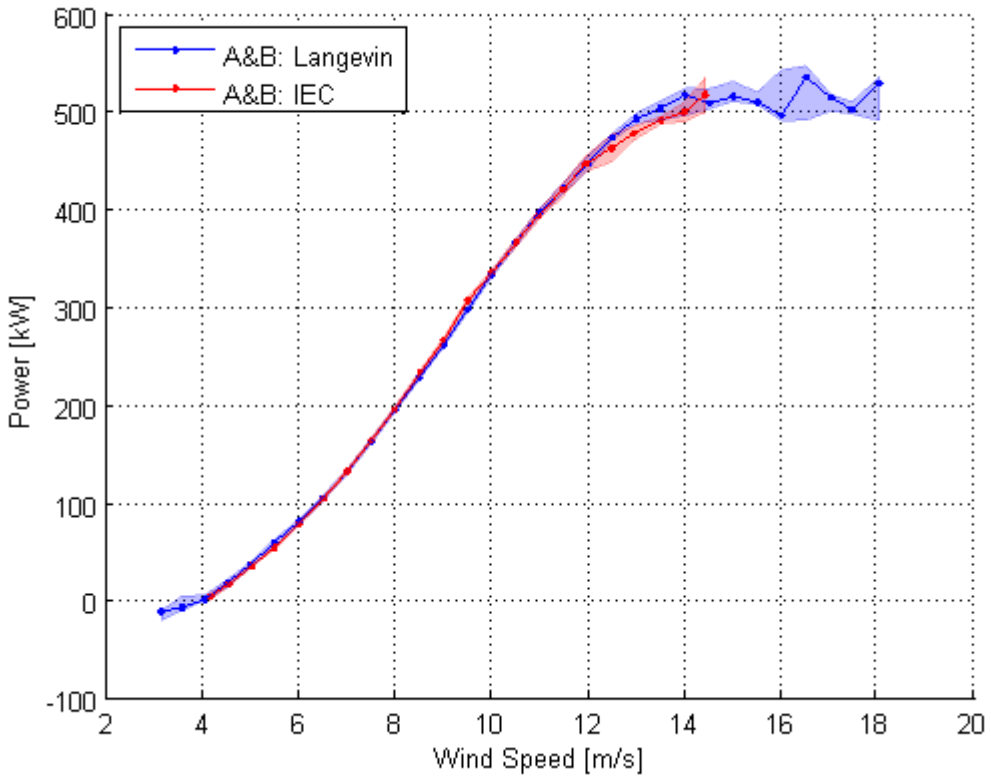

Figure 53 – Comparison of fixed point power curve from proposed revised Langevin method including pre-processing of wind speed with IEC power curve

As seen in Figure 53, data from both sectors were successfully combined to generate a sensible power curve which agrees with the IEC power curve. The data from sector B was successfully incorporated without throwing off the fixed points or introducing large amounts of error. If the datasets from sector B was combined without this pre-averaging, the results are severely affected; especially at low wind speeds.

Some method of pre-processing of the wind speed data may be required to help the Langevin method handle noise in the wind speed measurements. This method described above provides one potential solution to the problems discussed in this section related to sectors with low correlation between wind speed and power. A method for designing a filter for each wind direction sector could be developed to further investigate this method. It is suggested that the filter could be designed using the coherence plot for each sector in some way. Fully incorporating this sector based pre-processing of the wind speed into the Langevin method shows promise; it allows poorly correlated sectors to be utilized without ruining correlation of wind speed and power in well correlated sectors. This is one potential area for further work and has been recommended for further investigation in Section 8.

6 Investigation of Benefits of Langevin Method

Proponents of the Langevin method claim some benefits of the method over the standard IEC method. The validity of some of these claims was investigated in this section. The IEC method was used as the benchmark to compare against in these analyses.

The first claim investigated was that the Langevin method utilizes the measurement data more effectively since it can gain information from all the points in a 10 minute time series rather than reducing it to one point. Since it utilizes the data more effectively, the Langevin method should require less data to obtain the same level of uncertainty. This could allow the Langevin method to save time over the IEC method to obtain some desired level of convergence. The speed of convergence of the two methods was investigated, and the potential time savings were estimated as described in Section 6.1.

Proponents of the Langevin method also claim that it provides a way of estimating the fixed point power curve regardless of the turbulence level [9]. Since the Langevin method extracts its information form the power signal's response to fluctuations in wind speed, the magnitude of the fluctuations should not affect the fixed point predictions. The validity of this insensitivity to turbulence intensity of the Langevin method is discussed in Section 6.2.

6.1 Convergence of Langevin Fixed Points

Some of the potential benefits of the Langevin method rely on the assumption that the method can produce reliable results quicker than the IEC method. It is assumed that the information contained within the 10 minute time series which is not fully utilized by the IEC method will allow the Langevin method to converge quicker. In this section, the validity of this assumption will be tested.

For this analysis, the data described the base case in Section 5.1 was used once again. In the base case, additional filtering was done to reduce the measurement sector to wind directions between 260° and 310° based on the 10 minute mean wind direction. This was done so that pre-averaging, as described in Section 5.4, was not required. The Langevin inputs were also taken from the base case.

6.1.1 Measure of Convergence of Fixed Points

Over time, a given wind speed bin will accumulate data and as the amount of data within a bin increases, the uncertainty of the fixed point estimate should decrease. Estimates of the fixed point at different points in time should fluctuate but eventually converge upon some value. For example, a wind speed bin centered on 9 m/s with large amounts of data was used. Estimates of the fixed point at different points in time through the measurement period are shown in Figure 54. The fixed point power production was estimated using both the Langevin and IEC methods. The amount of data used for each estimate increases over time as shown with the sum of time spent in the bin.

As seen in Figure 54, the sum of time spent in a bin is not the same for the IEC and Langevin methods. This is because the data is divided between the bins slightly differently. A 10 minute mean wind speed may fall into one bin contributing 600 seconds of cumulative time in that bin for the IEC method; while, for the same 10 minute time series, using the Langevin method, the wind speed can fluctuate between the adjacent bins contributing some part of the 600 seconds to each of the bins it dwells in. This causes the difference between the "sum of time in bin" plot between the IEC and Langevin methods.

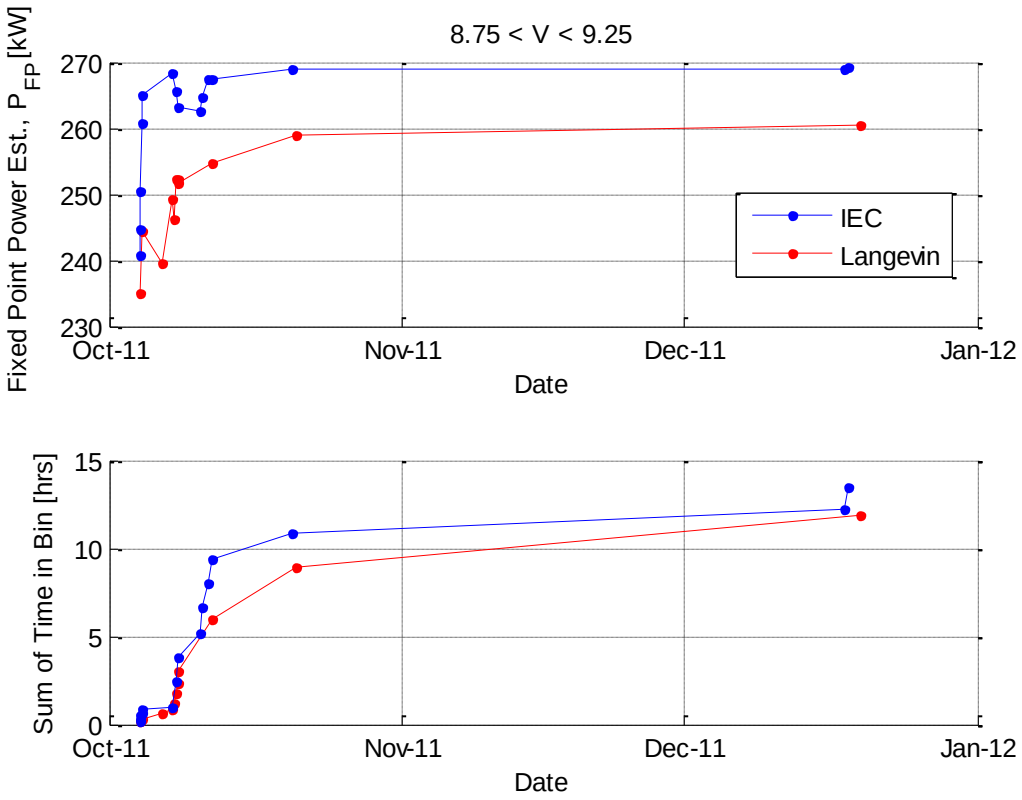

Figure 54 – Estimates of fixed point power production at different points in time using increasing amounts of data

As seen in Figure 54, the fixed point estimates vary over time but the fluctuations of the estimate decrease as the amount of data in the bin grows and the estimate converges to its final value. To show that the Langevin method is converging faster than the IEC method, some method of estimating the convergence is required.

Using the data chronologically, as shown in Figure 54, provides one estimate of the convergence of the fixed point estimates. To obtain more estimates of convergence, another resampling method was developed. Similar to the "bootstrap time series" method described in Section 4.3.3, segments of the time series were taken at random and reassembled into a new bootstrap time series. Time series segments were added until the number of points in a given bin was equal to the original chronological time series. With a bootstrap time series, the fixed point power of the bin could then be estimated using different amounts of the data. With many bootstrap time series, the probability distributions of estimates for a given amount of data were used to give a 90% confidence interval of the estimate.

An example of the convergence of both the IEC and Langevin methods is shown in Figure 55. For the two methods, segments of 1200 seconds were used when constructing the bootstrap time series. This period was chosen based on an estimate of the time taken for the wind speed signal to decorrelate. With the Langevin method 200 bootstraps were used to estimate the confidence interval. With the IEC method, since it is computationally much faster than the Langevin method, 1000 bootstraps were used to estimate the confidence interval. A 90% confidence interval was used for both methods.

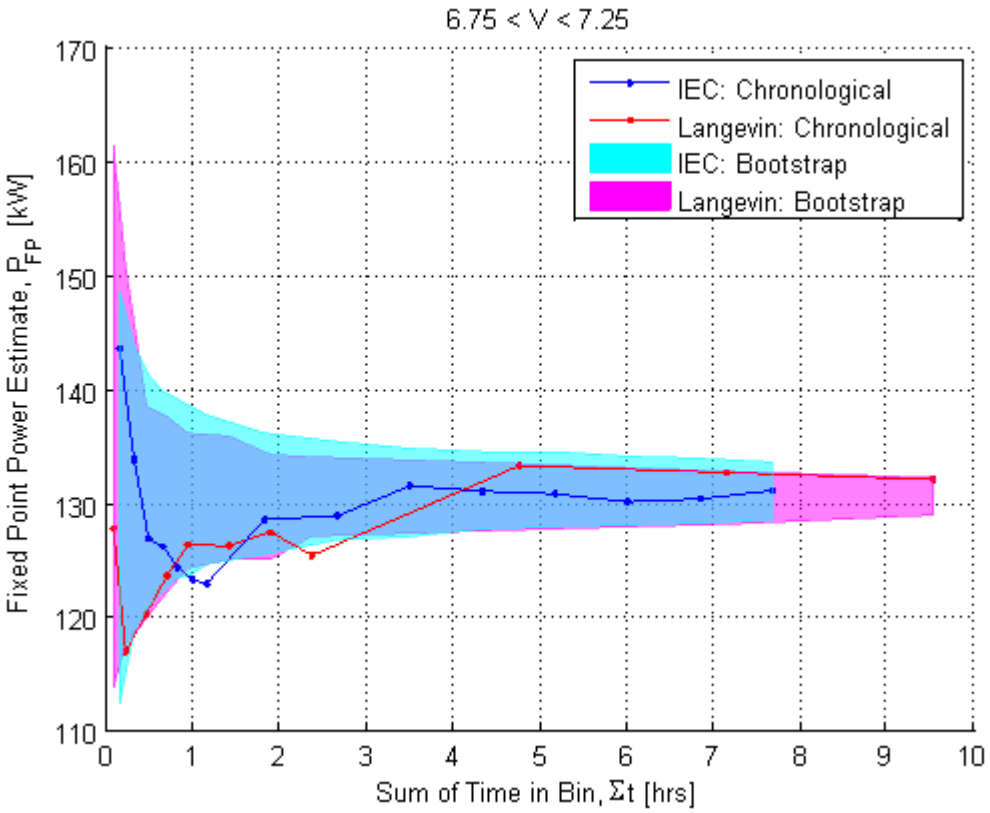

Figure 55 – Convergence of fixed point for one wind speed bin centered on 7 m/s

As seen in Figure 55, the variance of fixed point predictions decreases for both methods as the amount of data used increases. This plot helps indicate how much fixed point estimates can vary with a given amount of data. The plot can be used to estimate how much data is required to obtain a certain level of convergence. The speeds of convergence of the two methods are investigated in Section 6.1.2.

Additionally, Figure 55 shows how much uncertainty can exist with the minimum bin population required by the IEC standard of three 10 minute averages in each bin. In this bin shown in the figure, the range of error is approximately 20°kW when only three 10 minute mean data points are used (sum of time in bin of 0.5 hrs). In high wind speed bins, it can be difficult to obtain more than three data points since the probability of these high wind speeds is so low. If the three 10 minute mean power measurements have similar values, the uncertainty of the bin could be significantly underestimated. Since the high wind speed events are so unlikely, it is possible that the three measurements could be taken from the same extreme weather event possibly resulting in similar mean power measurements and an underestimation of the bin uncertainty.

As seen in Figure 55, the confidence intervals are plotted against the sum of time in the wind speed bin. This is similar to plotting against the “number of data points”; however, in the IEC method, one data point corresponds to 600 seconds while in the Langevin method, one data point is only 1 second. This axis was used instead of “number of data points” so that the two methods could be more easily compared.

6.1.2 Comparison of Speed of Convergence of Langevin and IEC Methods

The number of data points required to obtain some level of convergence can be used to show the speed of convergence of the two methods. The range of the confidence intervals shown in Figure 55 can be plotted as shown in Figure 56. To determine how much data is required to obtain a given uncertainty, the desired uncertainty can be chosen and the amount of time can be interpolated from

the graph. For example, the amount of data required to obtain an uncertainty range of 10 kW is shown in the figure.

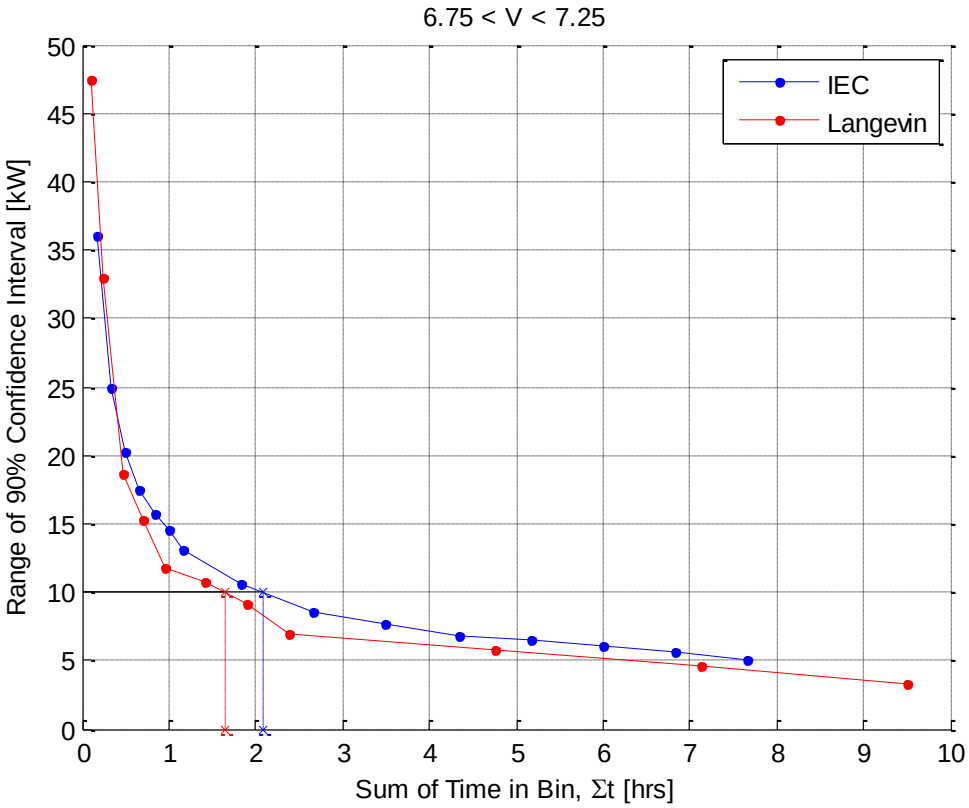

Figure 56 – Range of confidence intervals with increasing amount of data for one wind speed bin centered on 7 m/s

This process can be repeated for all wind speed bins to estimate how much data is required to obtain a desired level of uncertainty. An estimate of the amount of data required for each wind speed bin to obtain convergence of 10 kW and 20 kW is shown in Figure 57. The amount of data required for the IEC and Langevin methods is compared.

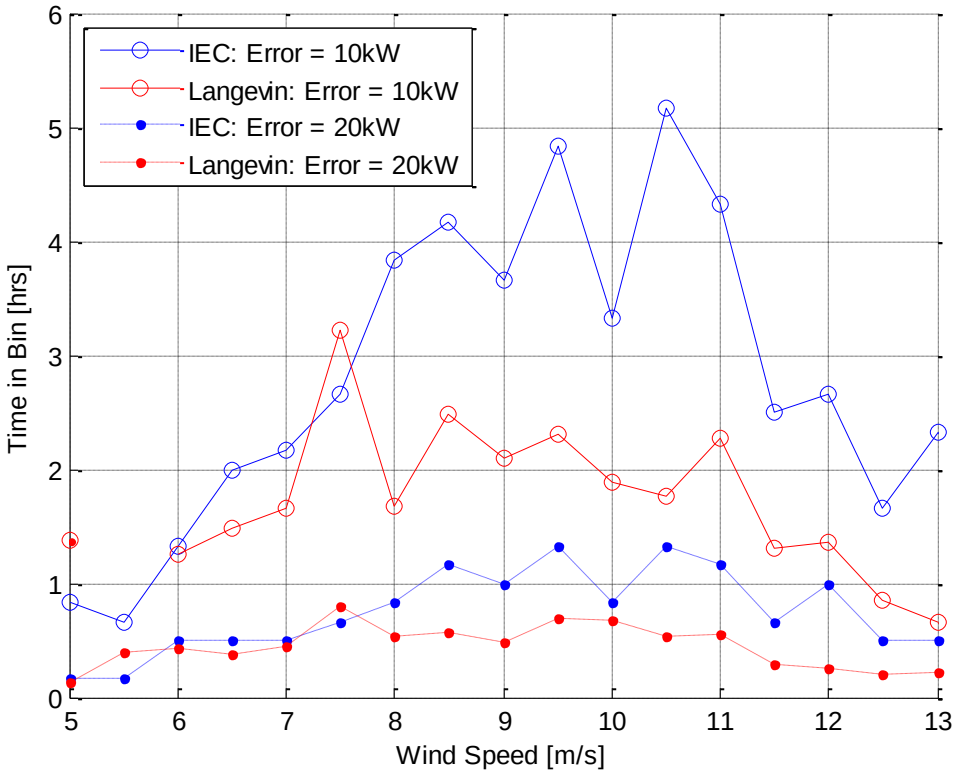

Figure 57 – Comparison of estimates for IEC and Langevin methods of time required in each bin to obtain desired levels of convergence of 10 kW and 20 kW

As seen in Figure 57, in a majority of the wind speed bins, the Langevin method is predicted to require less data than the IEC method to obtain the same level of uncertainty. This shows that the Langevin method may live up to its claim of reducing the amount of time required to obtain a reliable power curve. The amount of time savings is estimated and discussed in Section 6.1.3.

6.1.3 Length of Time Required to Measure Power Curve

In the previous section, it was shown that the Langevin method may require less data than the standard IEC method to obtain a power curve with the same level of uncertainty. The amount of time that this reduction in the required data can save when generating a power curve was estimated.

To estimate how long each bin will take to fill, wind speed data for a given test site can be simulated. Using simulated data allows the length of time required to fill a given bin to be estimated many times giving an estimate of the mean time period and distribution of time periods required to fill the bin. Additionally, using simulated data allows the fluctuations in wind speed and seasonal variability of the wind conditions to be included in the time period estimate.

In this report, a simple first estimate of the amount of time required to fill a bin was done using only the Weibull distribution and assuming each wind speed measurement is independent from the adjacent measurements. This does not reflect the true nature of wind speed measurements since it neglects the autocorrelation of the wind speed; however, it can be used to determine a lower bound on the estimate of the measurement time. Since true wind speed measurements are not independent, the actual measurement time will likely be higher than what is estimated using this method. Additionally, the IEC method may be penalized more than the Langevin since, with 10 minute averages, there will be fewer events possibly requiring a longer period of measurements for the distribution of wind speeds to fit to the Weibull distribution.

For an example, the best fit Weibull distribution for the site at Risø was estimated as shown in Figure 58. The 10 minute mean wind speeds at hub height from the met mast upstream of the Nordtank turbine were used to estimate the Weibull distribution.

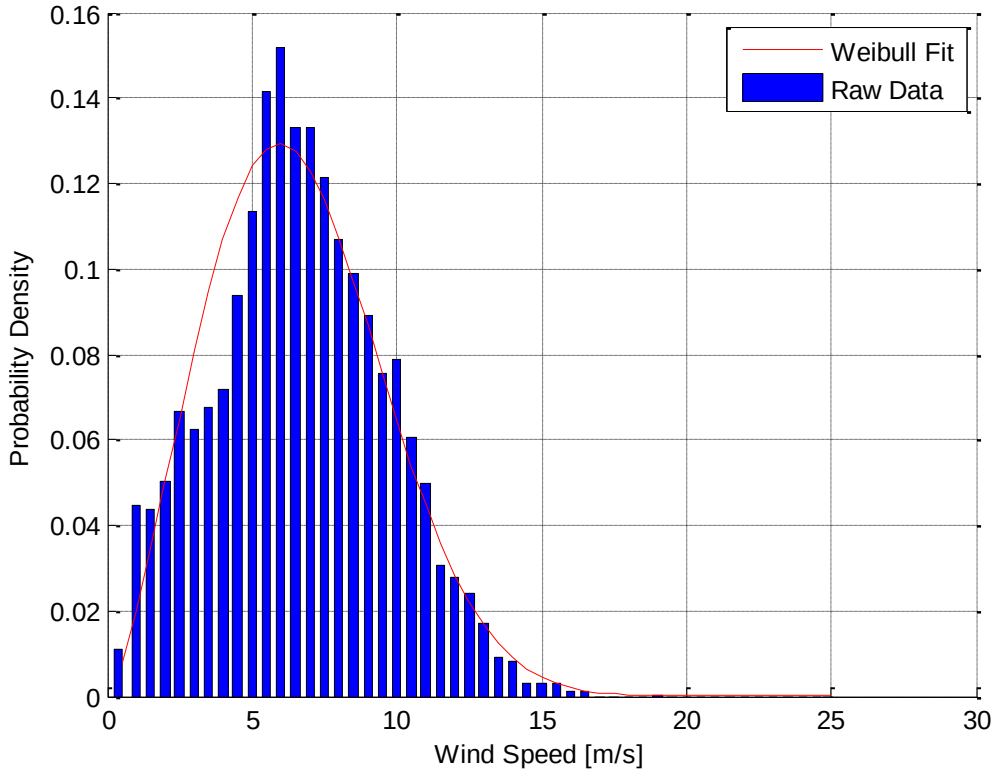

Figure 58 – Best fit Weibull distribution based on 10 minute mean hub height wind speeds for met mast at Risø site

The scale and shape parameters for the best fit Weibull distribution shown above are 7.49 m/s and 2.37, respectively. This Weibull distribution was used to estimate the average time it would take to fill each of the bins to obtain uncertainties of 10 kW and 20 kW as shown in Figure 59. The time estimate assumes the wind speed measurements were independent, the turbine is 100% available, and all the measurements can be used (e.g. the wind direction is never such that the met mast is in the turbine's wake).

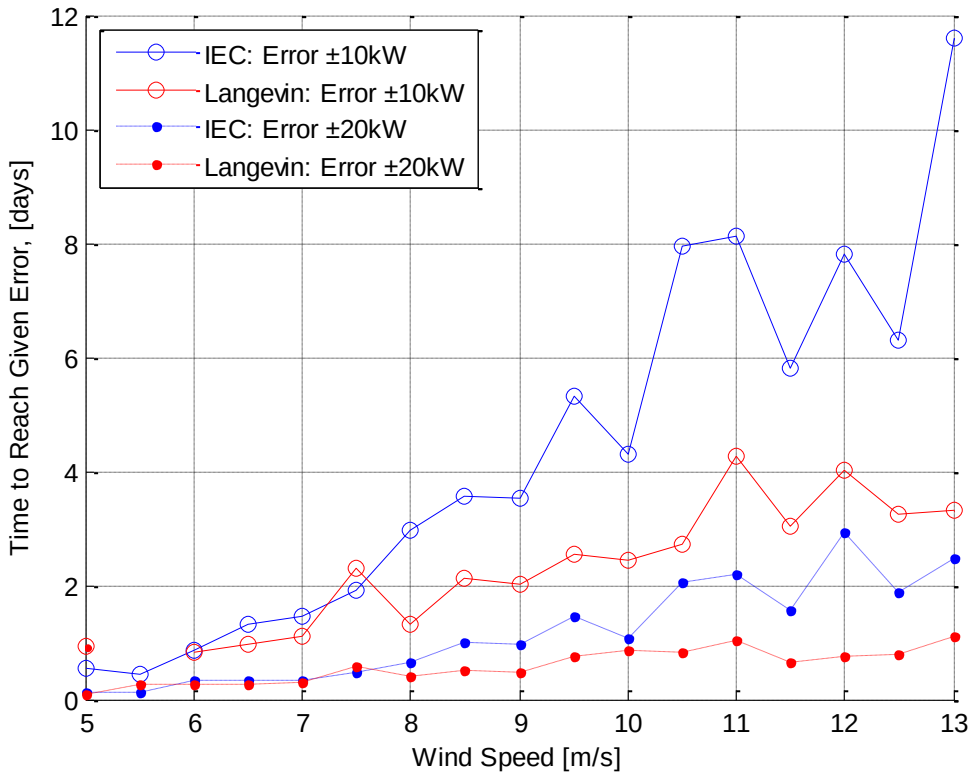

Figure 59 – Estimate of average time required to fill each wind speed bin at the Risø site to obtain uncertainties of 10 kW and 20 kW

As shown in Figure 59, in general, the Langevin method is estimated to require less time than the IEC method. The benefit of using the Langevin method is not apparent at low wind speeds where data is abundant; however, in the sparsely populated bins at high wind speeds, the reduction in the required data significantly lowers the measurement time required. Using the Langevin method, some of the high wind speed bins are predicted to converge to within 10 kW several days before the IEC method. Additionally, since this was a relatively simple initial analysis, the time savings could be even greater. If the turbine is not available 100% of the time, or if some of the measurements are un-usable for some reason the required time will be higher. Also, if the wind speed measurements were not assumed to be independent, the required measurement time would likely also be longer; however, this is expected to affect the IEC method more, potentially causing the benefits of the Langevin method to be even more exaggerated.

During power curve validation using the IEC method, measurements all the way up to 1.5 times the wind speed at 85% of the rated power are required [1]. For example, for the Nordtank turbine, the last bin would be centered on 17 m/s. There was insufficient data to estimate the convergence at this high wind speed but it is assumed that the Langevin method requires less data than the IEC method, as in most of the other wind speed bins. Additionally, for the Langevin method, this data requirement may be able to be fulfilled by short duration high wind speed transients while the IEC method needs to wait for three 10 minute periods of sustained high wind speeds which are quite rare. Power curve validation using the IEC method can be delayed while waiting for these high wind speeds. If the Langevin method continues to be developed it could eventually become accepted as an alternative method for power curve validation. The time savings provided by the Langevin method for these high wind speed bins may be even greater during validation.

These results show that there may be some truth to the belief that the Langevin method could be used to reduce measurement times. The Langevin method is predicted to converge faster potentially

alowing the power curve to be quickly determined with more certainty than the IEC method. This could help provide more certainty in the fixed points after a short duration of measurements or could help significantly reduce the measurement time during full power curve validation.

6.2 Turbulence Insensitivity of Langevin Method

One of the claimed benefits of this Langevin method is that it is insensitive to the level of turbulence so it should be able to extract the same fixed point power curve regardless of what level of turbulence the turbine is exposed to. The validity of this claim was investigated in this section. The performance of the Langevin method in high and low turbulence was compared against the performance of the IEC method which is known to be sensitive to turbulence as described in Section 2.3.1.

Once again, the data described in Section 3 was used. Additional filtering was done to reduce the measurement sector to wind directions between 260° and 310° based on the 10 minute mean wind direction at hub height. This was done so that pre-averaging, as described in Section 5.4, was not required. This filtered data set was then split into two data sets based on a data point's turbulence intensity. The median turbulence intensity of the data set was determined to be 11.5% and this was used to split the data set into two equal parts.

The Langevin method was then applied to each of the two data sets. The inputs of the Langevin method were the same as the base case described in Section 5.1. The resulting Langevin power curves from the two data sets are shown in Figure 60. The difference between the two power curves with the low turbulence power curve as the reference is shown in Figure 61.

Using the same two data sets, the IEC method was applied. The IEC power curves from the two datasets are shown in Figure 62. The difference between the two IEC power curves is shown in Figure 63, once again, with the low turbulence power curve as the reference.

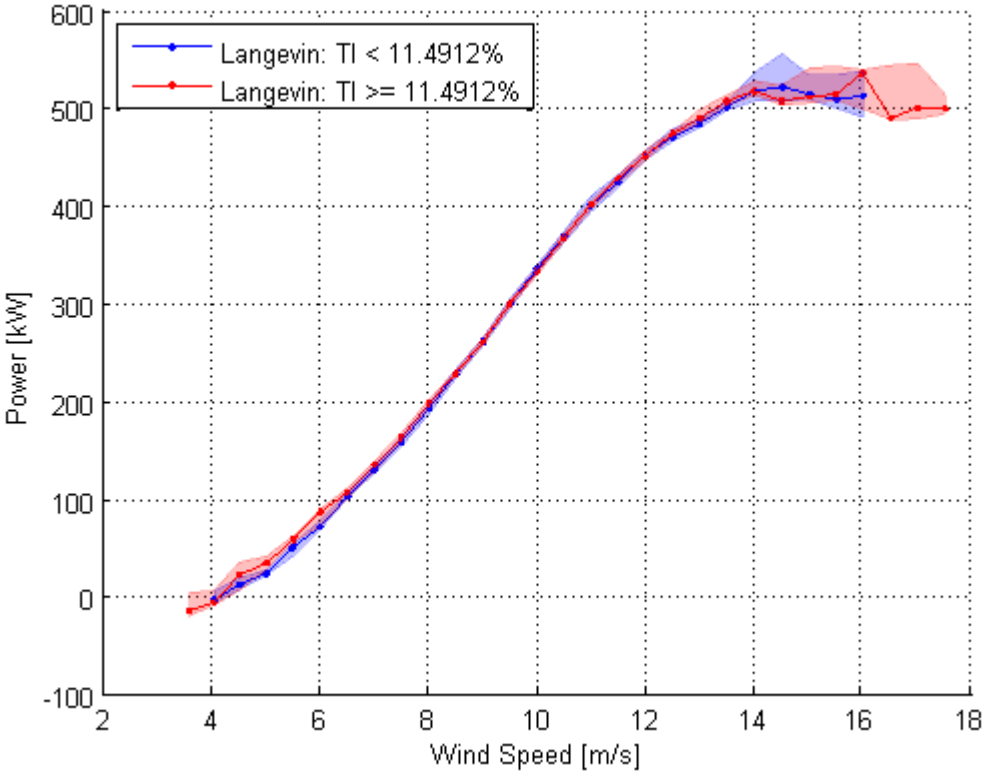

Figure 60 –Langevin fixed point power curves for high and low turbulence intensities

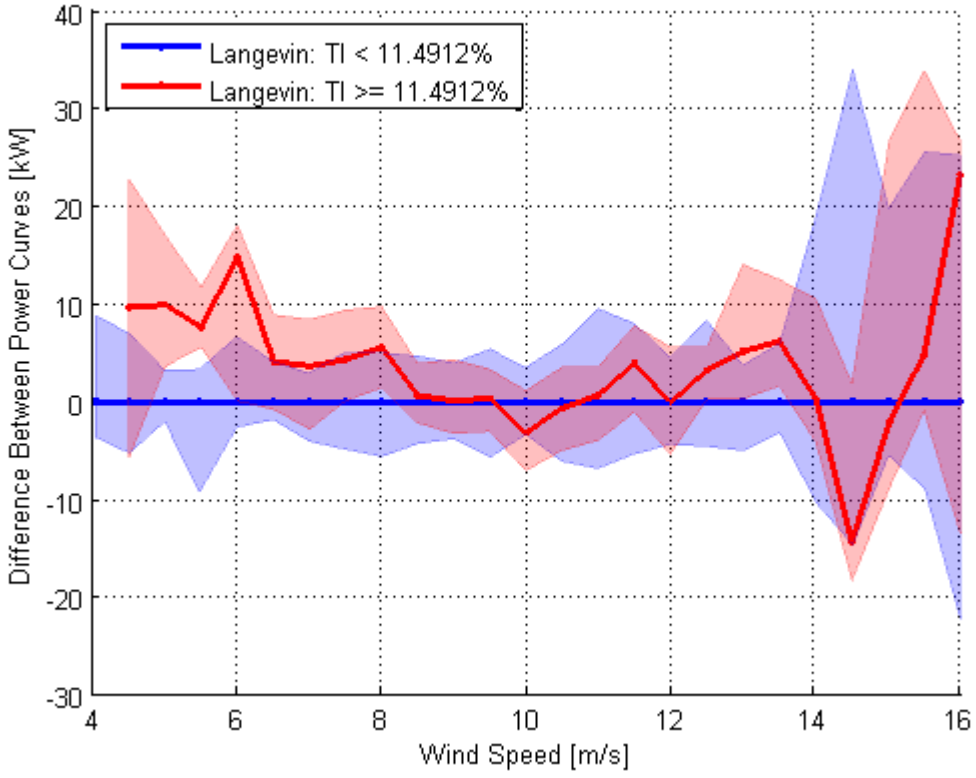

Figure 61 – Difference between Langevin fixed point power curves for high and low turbulence intensities

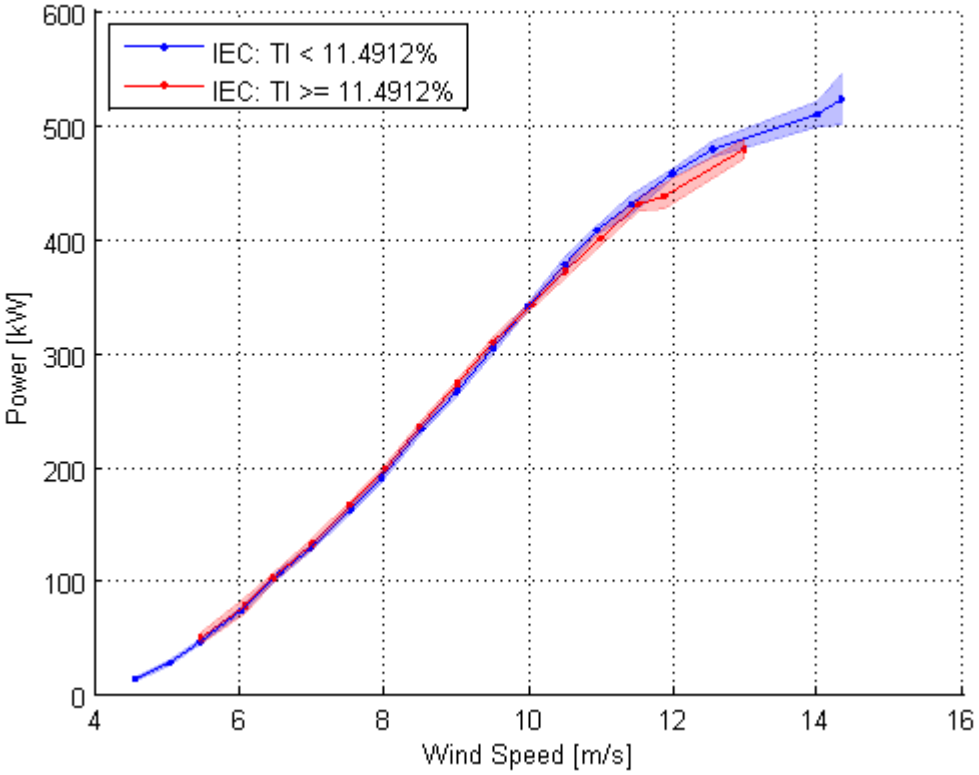

Figure 62 – IEC power curves for high and low turbulence intensities

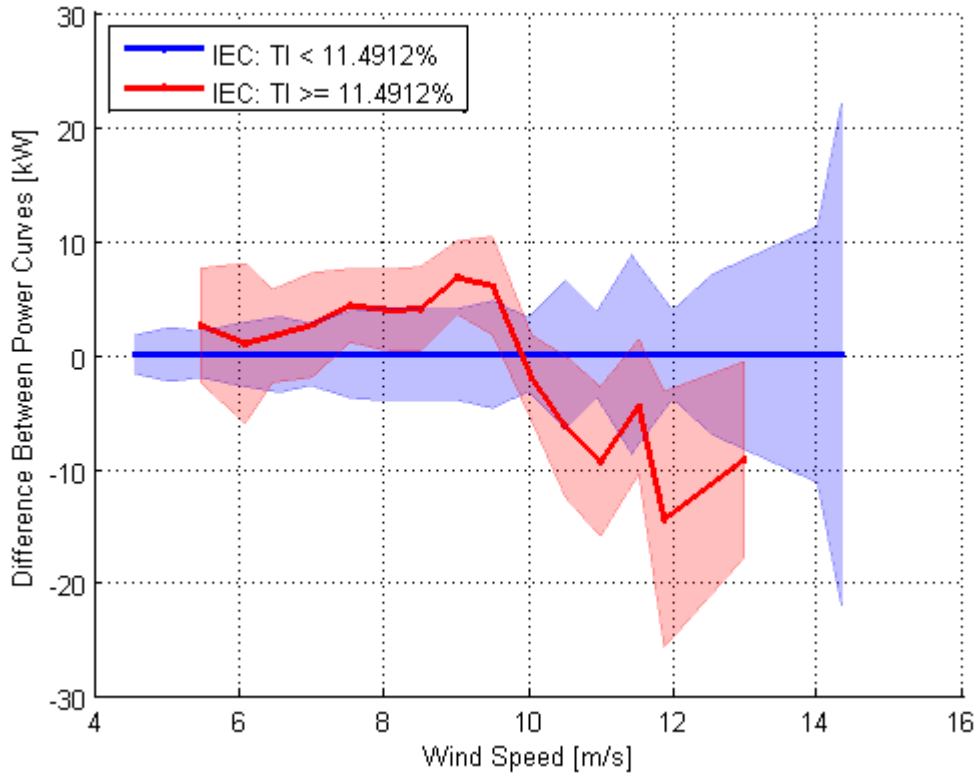

Figure 63 – Difference between IEC power curves for high and low turbulence intensities

As seen in Figure 60 and Figure 61, the two Langevin power curves agree and there is some overlap in the error bars for a majority of the power curve. The deviations at low and high wind speeds are likely due to a lack of data rather than the increase in turbulence intensity. The bin populations with the two data sets are shown in Figure 64. Furthermore, in the bins where both data sets have a large amount of data between about 7 m/s and 11 m/s, the Langevin method has predicted very similar fixed point estimates as shown in Figure 61. The small difference between the power curves in this range indicates that the Langevin method is converging to the same fixed points as the bins gain sufficient amounts of data.

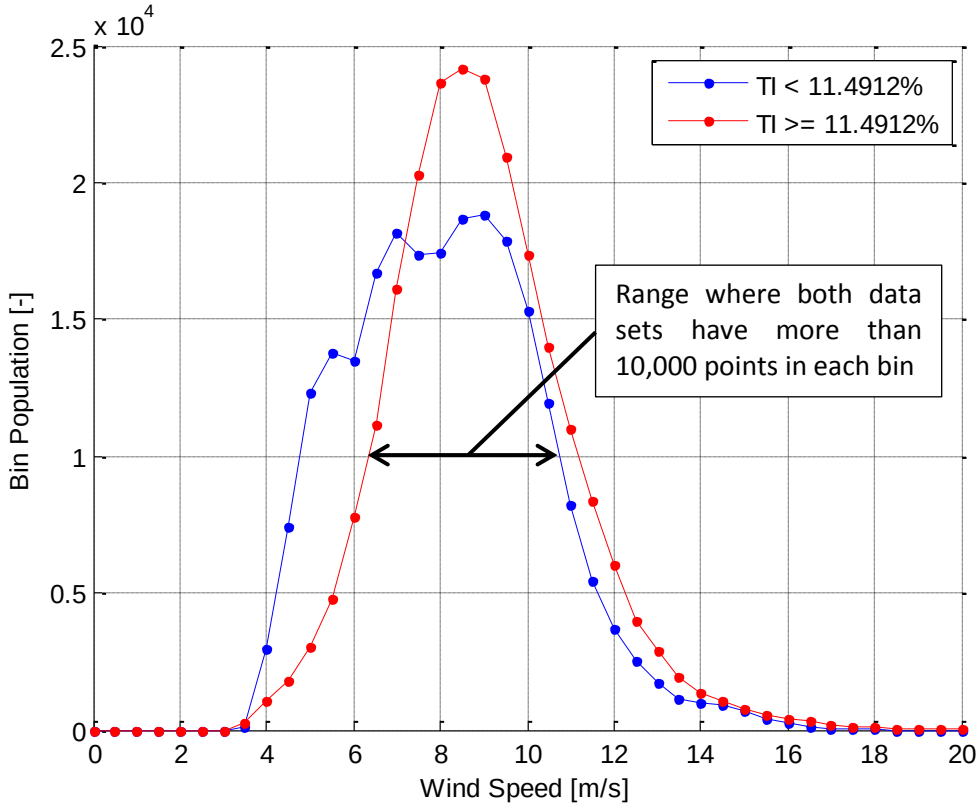

Figure 64 – Bin population for low and high turbulence data sets

The IEC method, on the other hand, is showing the same dependency on turbulence intensity that has been shown in the past and as discussed in Section 2.3.1. As seen in Figure 62 and Figure 63, the high turbulence power curve is predicting higher power values in the positive curvature region of the power curve below approximately 10 m/s. Above, 10 m/s the power curve inflects to a negative curvature and the high turbulence power curve is predicting lower power values.

The wind direction sector used in this analysis contains relatively low turbulence data since the wind is coming from in from the Roskilde fjord with low roughness and few obstacles. The turbulence intensity stayed below 15%, 90% of the time, and stayed below 19%, 99% of the time. The effects of turbulence on the IEC power curve are not as prominent as they could be if the turbulence level was higher. Additionally, since the data set was split in half, less data was available for convergence of the Langevin fixed points. In this analysis, the Langevin method has shown some promise that it may be less sensitive to turbulence than the IEC method; however, this should be tested with a larger data set with higher turbulence levels to see if the Langevin methods insensitivity to turbulence can be more thoroughly demonstrated.

7 Conclusions

This MSc. thesis performed at DTU Wind Energy as a part of the FastWind project has come up with the following major conclusions summarized in this section. The conclusions are discussed based in relation to the initial project goals.

The first goal of the project was to develop an understanding of and implement the Langevin method. A description of the Langevin method and how it was implemented in for this thesis project is included in Section 4. Some validation that this implementation of the Langevin method has been performed correctly and is providing reliable results was presented in Section 4.4. The source code containing the implementation of the Langevin method used in this project was made available for use and further development during the remainder of the FastWind project.

Secondly, the project objective was to test the practical application of the method on real measurement data. The necessary inputs and pre-processing required before applying the Langevin method were investigated to determine their effects on the results. It was found that some of the inputs to the Langevin method can significantly affect the fixed point predictions. Additionally, pre-processing of the wind speed data was shown to be necessary in some cases. With these conclusions, the Langevin method, as is, is not ready for industrial application and more development is required to improve the robustness of the results. A summary of the major conclusions that were found in the analyses of Section 5 are described as follows:

- The Langevin method can still function using only the hub height wind speed measurement. However, if more measurements are available, use of the shear correction is recommended.
- When performing fitting to the averaged response of the power signal for each sub-bin, the range of time lag chosen has a significant impact on the results obtained. When observing the averaged response of each sub-bin there is typically some approximately linear range before the response starts changing slope. This range is believed to depend on the speed of the dynamics of the turbine at each wind speed. In general, the low wind speed bins are found to have longer linear ranges than the high wind speed bins. Using the current Langevin method the range of time lag used for fitting is the same for all bins. Therefore, a range of time lag should be chosen such that, in the high wind speed bins, the fitting is not done using data which is far past the end of the linear region. Furthermore, enough data points should be used to obtain a reliable fit. For the Nordtank turbine, it was found that a range of 1 s to 10 s provides a reasonable fit in most of the wind speed bins. A more sophisticated method of dynamically choosing the range of time lags for fitting is described in the recommendations in Section 8.
- The Langevin method was found to be somewhat insensitive to the method chosen for binning the power data. To keep the uncertainty of the drift points approximately constant in all sub-bins, the method which dynamically sets the bin ranges to keep the number of data points in each sub-bin constant is recommended. This prevents large uncertainty at the ends of the drift curve.
- Some sort of pre-processing of wind speed before sending the data to the Langevin method may be necessary. If the small scale, high frequency, quickly evolving turbulence which does not correlate to fluctuations in the wind turbine's power production can be removed, the Langevin results are more reliable. This lack of correlation between certain frequency components of wind speed fluctuations was found to depend on wind direction. The uncorrelated frequency components could be removed using filters designed for each of the wind direction sectors. The pre-averaging of wind speed provides one method for filtering the wind speed data; however, a more sophisticated high order filter may perform better.

Recommendations for potential additions to the Langevin method are described in Section 8.

Finally, the last goal of the project was to investigate the validity of the claims that the Langevin method provides improved results over the IEC method. Using the Nordtank data, the claims that the method requires less measurement time than the IEC method was investigated. After these analyses, the Langevin method shows some indication that it may live up to the claims that are made by its proponents and development of the Langevin method may be worth pursuing further. A summary of the major conclusions found in Section 6 are described as follows:

- In general the Langevin method was found to typically require less data than the IEC method to obtain the same level of convergence. This gives some indication that the Langevin method may, indeed, provide some time savings over the IEC method when collecting data for power curve validation. This may be due to the fact that the Langevin method is taking better advantage of the data available to it by extracting more information from each 10 minute time series. This allows the Langevin method to predict the power curve with more certainty using the same measurement period.
- If the Langevin method were to someday be accepted as an alternative to the IEC method for validating wind turbine power curves, the time savings could be even greater. Since the power curve must be validated up to high wind speeds which are relatively infrequent, the power curve validation must wait for sufficient data to fill these high wind speed bins. The Langevin method can take advantage of short term wind speed dynamics while the IEC method requires three 10 minute periods with high average wind speed. The potential time savings for a full power curve validation is recommended for further investigation as described in Section 8.
- The Langevin method was found to be somewhat insensitive to turbulence intensity while the IEC method showed the same dependency on turbulence intensity that it has in the past. The wind speed bins which contained enough data for convergence predicted the same Langevin fixed points with both the high turbulence intensity and low turbulence intensity data sets. Recommendations to more thoroughly demonstrate the Langevin method's insensitivity to turbulence are described in Section 8.

This project has provided some initial results for the FastWind project and is a first step towards fulfilling the project goals. This project has shown that the Langevin method has some challenges associated with its implementation; however, these challenges are not insurmountable. Furthermore, the project has shown that the claimed potential benefits of the Langevin method may have some truth to them and the Langevin method is worth further investigation. To continue towards fulfilling the goals of the FastWind project, work to overcome the challenges highlighted by this thesis project and better understand the potential benefits of the Langevin method should be continued. Based on the conclusions of this report, some recommendations for further work have been made as described in Section 8.

8 Recommendations for Further Work

After the completion of this thesis project which contains many preliminary analyses on the Langevin method, some recommendations on how to proceed with the FastWind project have been made. These recommendations are suggestions to help improve the robustness of Langevin method or prove the potential benefits of the method. These suggestions should help the FastWind project develop a reliable method which can be promoted for use in the Danish wind energy industry. The recommendations are summarized as follows:

- The influence of correcting the measurements for density was not investigated. If temperature and pressure measurements are available, the correction should be applied to see how much influence it has on the results.
- All of these conclusions and recommendations are based on the Nordtank turbine owned by DTU at the Risø site. This is a relatively old and simple turbine since it is a fixed speed, stall controlled turbine. Before progressing with any of these recommendations, it should be verified that the results obtained here are not unique to the Nordtank turbine. It should be checked that the tendencies of the Langevin method observed in this thesis can be seen on other more modern wind turbines as well.
- Using simulated data, further investigations should be performed. The reasons why the knee of the power curve cannot be accurately reproduced by the Langevin method should be investigated. Additionally, the averaged response of the power production from one sub-bin does not reproduce the theoretical response accurately and the root cause and potential solutions should be investigated.
- As discussed in Section 5.3.1, when performing fitting to the averaged response of the power signal for each sub-bin, there can be differences in the length of the linear range of the averaged response. Currently the same range of time lag is used for fitting in every bin, even though, in the low wind speed bins, the linear range is long while in the high wind speed bins the range is short. A different range of time lag could be chosen for each bin to ensure that as much data is used for the fit as possible while preventing flattening of the drift curve by fitting past the end of the linear range. Initially this can be done manually; however it may also be possible to develop an algorithm for choosing the range of time lag used for each bin. This would remove one input from the Langevin method and increase the automation of the method which may help in increasing the method's robustness. Currently the fit in each bin should be checked individually to ensure the range of time lag was appropriately chosen. Since there is typically some linear range of the response, the fitting range could theoretically use any part of this range and obtain similar results. To automate the selection of fitting range for each bin, the fitting range could be increased until the slope of the fit starts to significantly deviate, indicating the fit is starting to utilize points outside the linear range of the response. Once the limits of the linear range are found, the whole linear range should be used to estimate the drift of the bin. Using the full linear range uses the most possible points helping to average out errors while preventing the drift to be misestimated by using points from outside the linear range of the averaged response.
- As discussed in Section 5.4, removing the high frequency wind speed fluctuations which are uncorrelated with power fluctuations was shown to potentially improve the robustness of the results produced by the Langevin method. The pre-averaging method described in this thesis is one possible method but the use of a purpose built high order digital filter would provide more control over the gain on each frequency component. Alternative methods of filtering the wind speed should be investigated to see if improvements can be made over the pre-averaging method used in this thesis. Additionally, since the correlation of the high frequency wind speed fluctuations varies with wind direction, the method proposed in

Section 5.4.4 could be investigated further. By filtering wind speed measurements with a filter designed for each wind direction sector based on the coherence of the wind speed and power measurements for that sector, more measurement sectors, and therefore more data, may be used when performing the Langevin analysis possibly helping to reduce uncertainty or measurement time.

- As discussed in Section 6.1.3, the Weibull distribution was used to estimate the time required to fill each wind speed bin and the wind speed measurements were assumed to be independent of one another. The wind speed could be simulated allowing affects such as the autocorrelation of the wind speed to be included. This would allow a much more accurate estimate of the potential time savings using the Langevin method to be determined.
- As discussed in Section 6.1.3, during full power curve validation, the high wind speed bins are the most difficult to fill since they are typically rare weather events. The IEC method must wait for three 10 minute periods of sustained wind at these high wind speeds to meet the minimum requirements indicated by the standard. The Langevin method, however, may be able to take advantage of short gusts up to these high wind speeds. Since the method may not need to wait for these infrequent weather events, it may be able to save significant amounts of time when validating the wind turbine's full power curve. However, insufficient data was available for the Nordtank turbine to analyze the convergence of the fixed point predictions at these high wind speeds. A data set which was used for the validation of a full power curve using the IEC method could be used with the Langevin method. The convergence of the Langevin method could be investigated with this same data set; especially in these high wind speed bins to determine if time savings can be achieved over the IEC method.
- The Langevin method's insensitivity to turbulence intensity was shown to hold some truth; however, the data set used was small and had a relatively low level of turbulence. The analysis could be repeated for a more complete data set with a higher turbulence level to determine if, at higher levels of turbulence, the Langevin method can still converge to the same fixed points as with low turbulence data. Additionally, the benefits of the Langevin method over the IEC method regarding turbulence insensitivity may be more pronounced at higher levels of turbulence since the IEC power curve will deviate more as the turbulence level increases.
- One of the potential benefits of the Langevin method was not investigated in this thesis; its ability to be used for forecasting wind power production. The Langevin model which is fit to measurement data for the turbine can be used to estimate the turbine's power production under a fluctuating wind speed. If the next day wind speed and turbulence intensity can be predicted, the Langevin model can be used to estimate the power production from these wind conditions. This potential benefit was not investigated in this thesis even though it may be even more important than accurate prediction of the fixed point power curve. Since the main reason for generating a power curve is for forecasting future power production, this could change the way we think about power curves. If it can be shown that the Langevin model can more accurately predict the power production than the standard IEC power curve, then less emphasis can be put on determining the fixed point power curve. More work can be focused on being able to develop a model which can accurately predict the wind turbine's response to a fluctuating driving wind speed. The fixed points are a part of this model but they will no longer be the focus of the investigation as is currently the case.

9 Bibliography

1. **IEC.** 61400-12-1; Wind turbines - Part 12-1: Power performance measurement of electricity producing wind turbines. s.l. : International Electrotechnical Commission, 2005. IEC 61400-12-1.

2. *Markovian Power Curves for Wind Turbines.* **Anahua, Edgar, Barth, Stephan and Peinke, Joachim.** 2008, Wind Energy, pp. 219-232.

3. **Gottschall, Julia.** *Modelling the variability of complex systems by means of Langevin processes.* Oldenburg : University of Oldenburg, 2008.

4. **Hansen, Martin O. L.** *Aerodynamics of Wind Turbines.* London : Earthscan, 2008.

5. **Hansen, Svend O. and Dyrbye, Claës.** *Wind Loads on Structures.* West Sussex : John Wiley & Sons Ltd, 1997.

6. **Wagner, Rozenn.** *Accounting for the speed shear in wind turbine power performane measurement.* Lyngby : Technical Univerisy of Denmark, 2010.

7. *Characterization of the wind turbine power performance curve by stochastic modeling.* **Anahua, Edgar, Barth, Stephan and Peinke, Joachim.** 2007, Wind Energy, pp. 173-177.

8. *Turbulence Normalisation of Wind Power Curve Data.* **Albers, Axel.** Warsaw : s.n., 2010. Proceedings of the European Wind Energy Conference (EWEC) 2010.

9. *Power performance of wind energy converters characterized as stochastic process: applications of the Langevin power curve.* **Wächter, Matthias, et al.** 6, s.l. : Wiley, 2011, Wind Energy, Vol. 14, pp. 711-717.

10. *Stochastic Modeling of Wind Power Production.* **Milan, Patrick, Wächter, Matthias and Peinke, Joachim.** 2011. Proceedings of EWEA 2011.

11. *Quantitative Estimation of Drift and Diffusion Functions from Time Series Data.* **Kleinhans, D. and Fredrich, R.** 2007, Wind Energy, pp. 129-134.

12. **Wolter, Kirk M.** *Introduction to Variance Estimation.* s.l. : Springer, 2003.