Performance monitoring by tracking estimated power curves on a wind farm level

Nymfa Noppe
Offshore Wind Infrastructure Lab,
Vrije Universiteit Brussel

Abstract

A well-known parameter to detect performance issues is the availability (time-based or production-based). Since the availability is highly affected by down-time, small changes in production are usually not detected by the availabilities. To detect these small changes one can track the **power curve** over time. This is done by calculating the power curve on regular basis using data during **normal operation** and defining **health indices** related to the power curve. These health indices **compare the calculated power curve to the warranted power curve**, e.g. the area under or the distance between both power curves. By using these health indices a **seasonal change** in the turbine performance was observed and **underperforming** turbines are easily detected.

Introduction

An easy way to detect some performance changes is calculating the **availabilities** (time- and production-based). This is already done by a lot of operators. However the availabilities cannot capture all kinds of **underperformance** or **performance changes**.

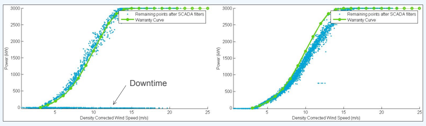

Figure 1: Different possible issues of wind turbine: low availability due to downtime (left figure), underperformance during normal operation with higher availability (right figure)

Figure 1 gives an example:

- based on the data set on the left: due to a lot of down-time, the availabilities will be low
- based on the data set on the right: quite normal (but not very high) production-based availability and a normal (high) time-based availability

Using the availabilities, the turbine on the left will be considered as being the one with issues. However, the turbine represented by the right turbine has some issues too: compared to the warranted power curve, it is slightly underperforming. To detect this kind of issue, another method is needed.

Objectives

The main objective is to find performance issues that remain unnoticed in the availability calculations. This can be done by regularly calculating the **power curve based on data of normal operation** of a wind turbine. By taking only data, based on normal operation, events such as down-time are removed and therefore cannot influence the resulting power curve. All other points, even during performance issues, remain part of the dataset.

So the main objectives are to:

- find a method to filter out data of normal operation and calculate the power curve regularly.
- define and use some metrics to compare those power curves to one-another and the warranted power curve.

Step 1 : Power Curve Estimation

To estimate the power curve of a wind turbine, a standard is already available as guideline (IEC 61400-12-2).

Following this standard, the needed steps are:

- first the data is filtered
- then several **corrections,** e.g. air density correction, are done
- finally a **power curve** is calculated from the remaining data

When applying IEC 61400-12-2 on the monthly data, quite often the suggested filters **reject nearly all data points** and no power curve can be calculated. This means the method should be adjusted to allow a regular calculation of power curves.

Therefore the filters suggested in the standard, among others, are integrated in the software of **DYNA***wind*. In **DYNA***wind* the user can define his **own combination of filters**. Some examples of the built-in filters are: use only data within the measurement sector, exclude blade icing, remove statistical outliers,... As such **less severe filters** can be applied on the data and enough data remains to calculate a power curve.

The last step is to calculate a **power curve** with the remaining data set. The method suggested by the standard is the **binning method**: divide the whole data set into bins based on the windspeed and calculate the mean value for the power of the datapoints of each bin. Figure 2 shows this technique applied on some random simulated data (without physical meaning): the red lines indicate the bin borders, the red dots the resulting bin values.

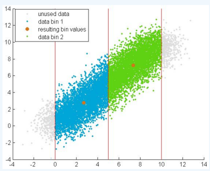

Figure 2: Example of binning

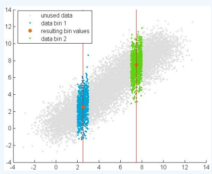

Figure 3: Example of k-nearest neighbour

DYNA*wind* also allows to calculate the power curve with a second algorithm: **the k-nearest neighbour technique**. Instead of dividing the dataset into several bins, a set of predefined values for windspeed are given. For each predefined value, the algorithm will search for the k data points with a value for windspeed closest to the predefined value and calculate the mean power value of those k data points. Figure 3 illustrates the working principle of this technique on the same simulated data: the red lines indicate the predefined values, the red dots are the resulting bin values.

Step 2: Monitoring the performance

To be able to compare curves to each other and the warranted curve, two metrics are defined. One is based on the **area under the curves**, another is based on the **distance** between the curves (figure 4).

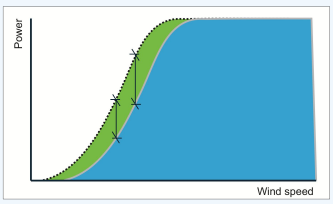

Figure 4: Health indices based on the power curve. The difference between the warranted curve (grey) and the measured power curve (dashed) can be defined by the distance between both curves (black) or by the ratio of the area under the measured power curve (green and blue) and the area under the warranted power curve (blue)

These two metrics can be used as health indices for the wind turbine. In the continuation of this poster, the health index based on the surface area under the curve is used.

Results…

Figure 5 gives the resulting values for the health index, based on the area under the curves, calculated on a monthly basis for a year. A **seasonal change** can be observed. The turbine represented by the blue line performs better than the others, but the curve is decreasing more than the others, which could indicate a performance issue of the turbine.

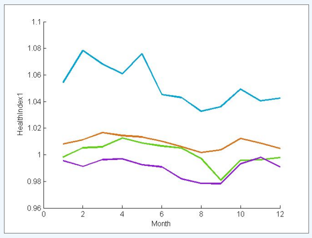

Figure 5: Health index for 4 turbines, monthly calculated for a year. A seasonal variation can be seen.

Figure 6 gives the resulting values for the health index, based on the area under the curves, calculated for a year for the complete wind farm. In this case the bigger the dot, the better the result for the turbine.

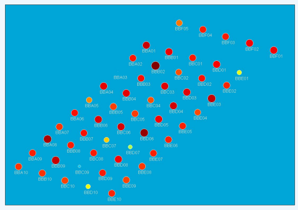

Figure 6: Health index calculated for 1 year, for all turbines of a wind farm. Turbines with performance issues are immediately detected as smaller dots.

Conclusions

By regularly calculating the power curve during normal operation and by tracking the changes in the health indices over time, even minor **changes in the production during normal operation** are easily detected. As illustrated, a seasonal change is one of them.

By comparing the results of the health indices of the turbines to each other, the **underperforming turbines** can easily be distinguished.

In the end, a method is given to detect small performance issues, which are not always visible using the availabilities.

References

[1] Andrew Kusiak et al. *On-line monitoring of power curves.* Renewable Energy 34 (1487-1493), 2009.
[2] International Electrotechnical Commission.*Wind Turbines - Part 12-2: Power performance of electricity producing wind turbines based on nacelle anemometry.* 2008 (Draft)

Acknowledgements

This work has been founded by the Institute for the Promotion of Innovation by Science and Technology in Flanders (IWT) in the framework of the OptiWind project: Serviceability of the next generation offshore wind turbines (www.owi-lab.be).