Importance sampling of severe wind gusts

René Bos[1], Wim Bierbooms, and Gerard van Bussel
Wind Energy Research Group, Delft University of Technology
Kluyverweg 1, 2629 HS Delft, The Netherlands
[1]r.bos-1@tudelft.nl

Keywords – extreme loads, wind gusts, Monte Carlo method, importance sampling

An important problem that arises during the design of wind turbines is estimating extreme loads with sufficient accuracy. This is especially difficult during iterative design phases when computational resources are scarce. Over the years, many methods have been proposed to extrapolate extreme load distributions from relatively short time series with "mean turbulence". In this work, however, we focus on finding the response to extreme gusts based on the ability to generate conditional turbulent wind fields. Load distributions can then be constructed on the basis of a Monte Carlo method with importance sampling.

I. THEORY

The most straightforward way to determine an extreme load distribution is by a crude Monte Carlo simulation. In this case, N ten-minute wind fields are generated from the mean wind speed distribution, $f(\bar{U})$, and fed to an aeroelastic model. When, for each sample, the maximum load is extracted, it results in a series of extreme loads $x_1, \dots, x_N$. An extreme load distribution then follows from

$$\hat{F}(L) \approx \frac{1}{N}\sum_{i=1}^{N} 1(x_i \leq L), \tag{1}$$

where

$$1(x \in S) = \begin{cases} 1 & \text{if } x \in S, \\ 0 & \text{if } x \notin S, \end{cases} \tag{2}$$

is the indicator function. The extreme load distribution is generally a cumulative distribution, representing the probability of non-exceedance. The return period then follows from

$$\mathrm{T} = \frac{1}{1-F}. \tag{3}$$

The return period should be interpreted as the time period after which the extreme value is exceeded once on average. For example, if extreme loads are extracted from ten-minute wind fields, the 90th percentile corresponds to a return period of 100 minutes. This means that, out of a sample size of 10, the highest load lies in the 90th percentile ($F = 0.90$) and can be called the 100-minute load. Perhaps unsurprisingly if judging by the name, the crude Monte Carlo method is not very effective; about $2.6 \cdot 10^6$ ten-minute wind fields are needed to reach the 50-year return level ($F = 0.9999996$).

In practice, 50-year loads are extrapolated from much smaller sample sizes because of the effort it takes to run aeroelastic simulations. This can be very difficult because the shape of the extreme load distribution can contain bends or curves that easily lead to bias [1]. Therefore, only a fraction of the data is really usable for fitting (say, the 5–10% highest loads) and a lot of computation time is needed to predict extreme loads with sufficient accuracy.

A common approach to reduce the uncertainty in Monte Carlo methods is to work with *importance sampling*. In this case, N samples are drawn from a particular distribution, $w(\mathbf{k})$, and are weighted by the *likelihood ratio*, $f(\mathbf{k})/w(\mathbf{k})$:

$$\hat{F}(L) \approx \frac{1}{N}\sum_{i=1}^{N} 1(x_i \leq L)\frac{f(\mathbf{k})}{w(\mathbf{k})}, \tag{4}$$

where $f(\mathbf{k})$ is the probability density function associated with the parameter space $\mathbf{k}$. After choosing a number of relevant parameters, the sampling distribution can be chosen such that the computational budget is efficiently spent on simulating severe events. Ultimately, this leads to much better predictions than what is obtained with a crude Monte Carlo method, where most of the extreme loads are cluttered around a mean.

Importance sampling becomes interesting when one has control over a large number of relevant parameters. In order to gain more control over the wind field, one can rely on the principle of *constrained stochastic simulation* [2,3]. This makes it possible to simulate a conditional turbulence field that adheres to a number of constraints. These constraints can be set such that a predefined extreme gust is embedded within the field (see Fig. 1). The wind field then follows from the mean wind speed, $\bar{U}$; the gust's amplitude, A; the gust's position, $\mathbf{x}_0 = [x_0, y_0, z_0]^\mathrm{T}$; the gust's length scales, ℓ_x, ℓ_y, and ℓ_z; and the gust shape (i.e., rectangular, ellipsoidal, etc.). Based on random field theory, it is possible to find the probability associated with such events [4]. What remains is finding out which combination of parameters lead to the most severe load cases.

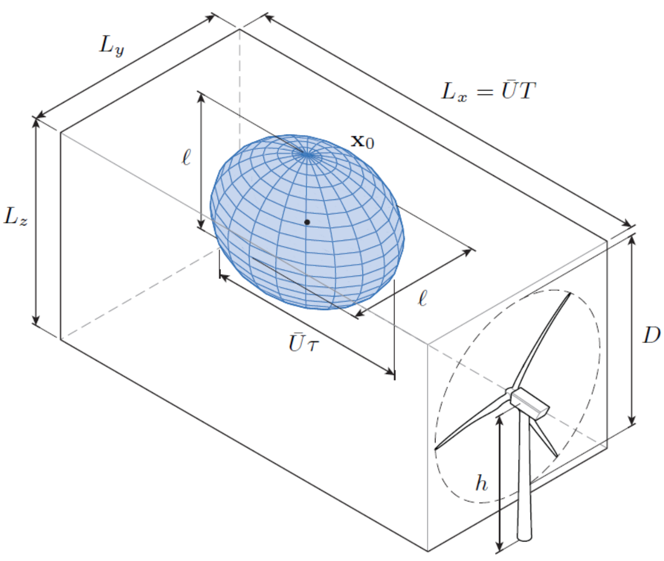

Fig. 1 Sketch of a gust in a constrained wind field, showing all the relevant parameters.

II. Methodology

A. Reference data set and model set-up

In this paper, an importance sampling method will be compared to a crude Monte Carlo simulation of 96 years of operation [5]. The exact same model set-up is used, but the wind fields are, of course, generated outside of TurbSim.

Load calculations were performed on the onshore version of the NREL 5 MW reference turbine using FAST v7. The wind climate is modeled according to an IEC class 1B site[1] for which the ten-minute mean wind speed follows a Rayleigh distribution. Each simulation was run for 2 minutes, where the first minute was discarded to avoid any start-up transients. As with the reference data set, turbulence was generated according to the IEC Kaimal spectrum [6] on a 20 x 20 grid with a width and height of 137 m and a temporal frequency of 20 Hz.

The turbulence grid turned out to be quite coarse in the y- and z-directions. This meant that averaging the amplitudes out over a volume was not was not as effective as hoped. Moreover, there is also no clear relationship between the position of the gust and the load (something that does often exist for other turbines). Therefore, the choice was made to stick with single-point gusts for this exercise.

Another consequence of the coarseness of the grid is that analytical approximations for the gust probability (e.g., as explained in [4]) lose their validity as they require a smooth random field. This meant that the probability of gusts occurring had to be derived empirically. A generalized extreme value distribution was therefore fitted to the velocity maxima found in 13,000 ten-minute wind fields for each wind speed between 3 and 25 m/s, yielding a joint probability density function $f(\bar{U}, A)$. From this, it is also possible to derive the 50-year return level for the gust amplitude (plotted as a dashed line in Fig. 3), which seems to be at around 6.0σ. At higher wind speeds, the 50–year gust is found at a slightly higher amplitude, because a higher rate of advection means more gusts are counted in the same time period.

B. Response to extreme gusts

Based on the reference data set, we can already see at which wind speeds the highest loads are found (see Fig. 2). Usually, we would expect to see the highest loads around the rated wind speed (11.4 m/s). This makes sense, seeing as the rotor operates at a maximum thrust at this point. However, the actual 50-year extremes are found at much higher wind speeds.

In order to investigate this further, and to expose dependencies of the extreme loads on the gust amplitude, 50,000 point-gusts (i.e., gusts with zero volume) were uniformly sampled from $\bar{U} \in [3, 25]$ m/s and $A/\sigma \in [-10, 10]$. The responses were binned and, for each bin, the maximum load is plotted in Fig. 3. This shows that for positive amplitudes, the results are indeed as expected. The high loads found at beyond-rated wind speeds actually correspond to negative amplitudes (i.e., a sudden drop in wind speed). This might seem counter-intuitive at first and is best explained by showing an example of a time series extracted from a set of 10-minute load cases (see Fig. 4). In situations like these, it appears that the pitch controller is unable to handle extreme drops in wind speed. Such behavior is normally dealt with by an additional nonlinear gust controller, but this is not included in the baseline controller of the NREL 5 MW turbine [7]. The same kind of events are also noted in literature [1,5].

C. Sampling distributions

From Fig. 3, it is fairly straightforward to derive the conditions at which the extreme blade root bending moments can reasonably be found. For this exercise, it is assumed that the extreme loads are dependent on 2 parameters, namely the mean wind speed, $\bar{U}$, and the gust amplitude, A. The gust's position, (y_0, z_0), is assumed to be uniformly distributed over the yz-plane. The gust's time stamp, $x_0/\bar{U}$, is also varied uniformly such that the rotor's azimuth angle can be anywhere between 0 and 120° at the time of impact. A sampling

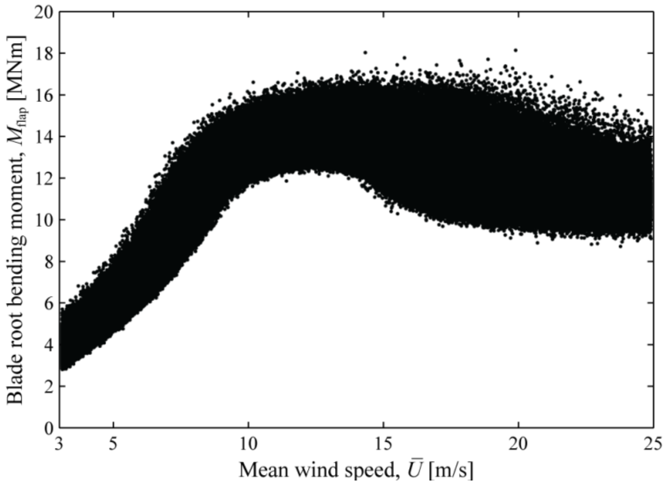

Fig. 2 Extreme blade root flapwise bending moments obtained from the reference data set, plotted as a function of the mean wind speed $\bar{U}$.

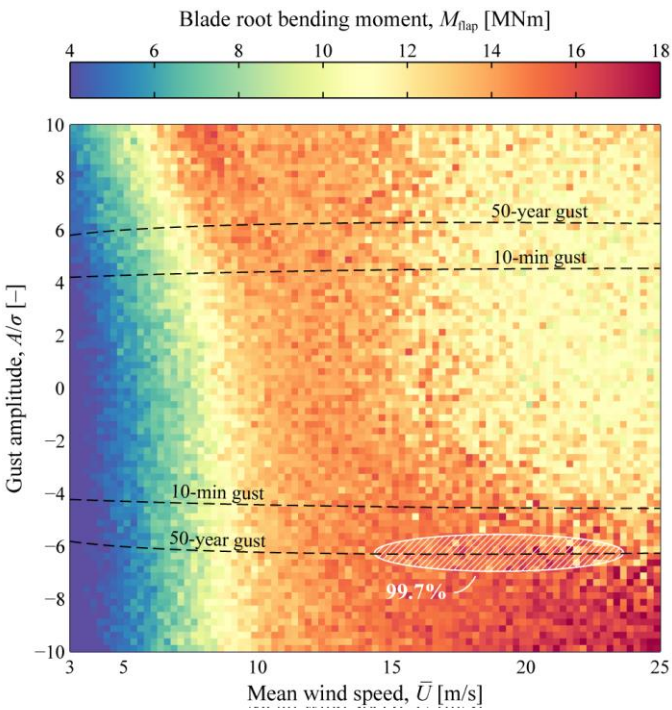

Fig. 3 Extreme blade root flapwise bending moments in response to a gust of amplitude A at a mean wind speed $\bar{U}$. The white hatched area marks 99.7% (±3σ) of the sample space.

[1] The original paper [5] specifies a class 2B site, but this has been corrected with the release of the data set (see http://energy.sandia.gov/?page_id=13173).

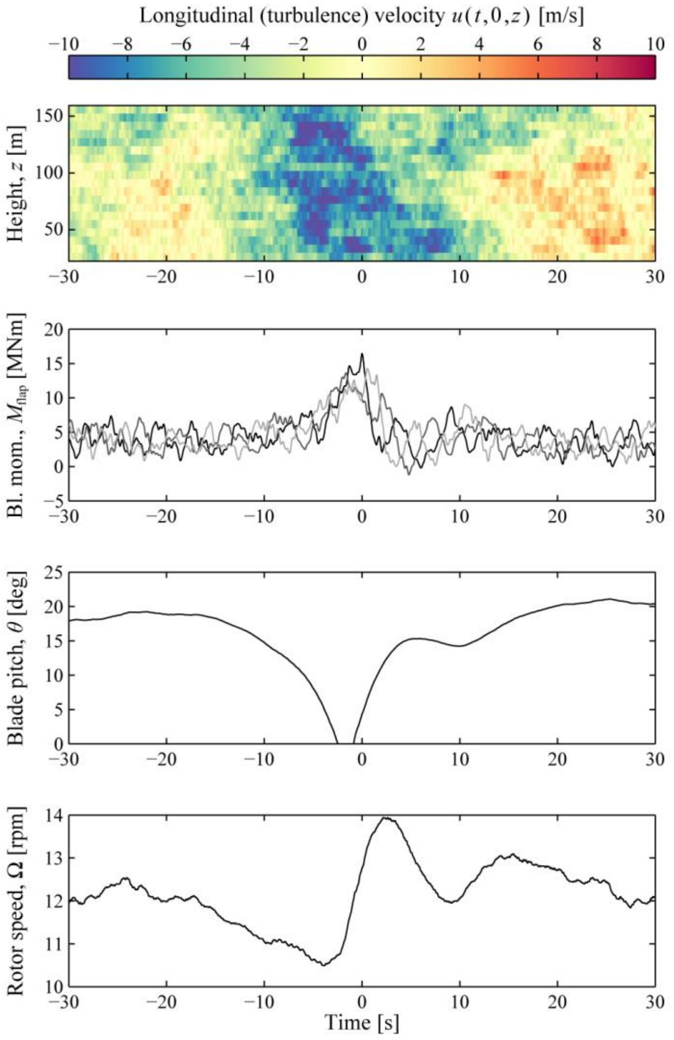

Fig. 4 An extract of a ten-minute time series at $\bar{U}$ = 22 m/s with an extreme blade root bending moment of 16.4 MNm.

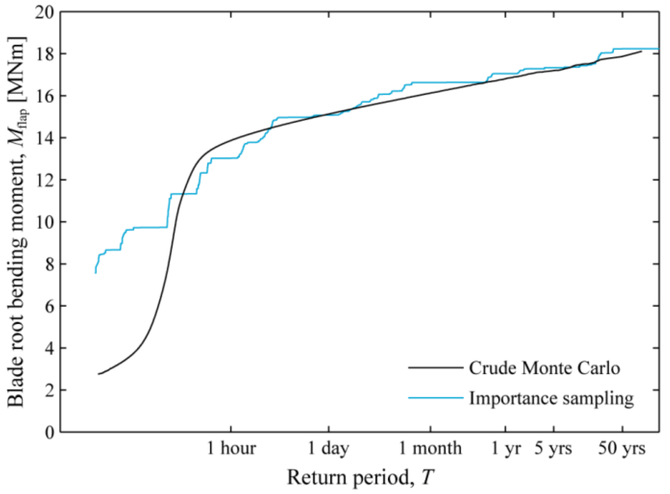

Fig. 5 Return level plot of the extreme blade root flapwise bending moments, constructed from a crude Monte Carlo simulation ($N = 5 \cdot 10^6$) and an importance sampling method ($N = 1 \cdot 10^5$).

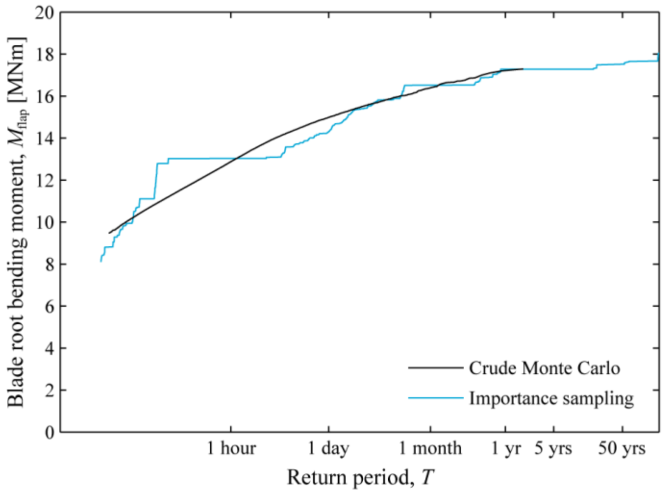

Fig. 6 Return level plot of the extreme blade root flapwise bending moments at 18.5≤ $\bar{U}$ ≤ 19.5 m/s, constructed from a crude Monte Carlo simulation ($N = 1 \cdot 10^5$) and an importance sampling method ($N = 3 \cdot 10^3$).

distribution is defined as a multivariate normal distribution, $w(\bar{U}, A)$, with a mean at $\bar{U} = 19$ m/s, $A = -6.25\sigma$ (see Fig. 2). This is to concentrate the computational effort on events that occur several times during a turbine's lifetime, as well as on very rare events. In any way, it is clear from Fig. 2 and 3 that we need to be looking beyond the rated wind speed.

III. Results

Fig. 5 shows the extreme load distributions that arise from a crude Monte Carlo simulation ($N = 5 \cdot 10^6$) and the importance sampling method ($N = 1 \cdot 10^5$). The distributions are plotted on a double logarithmic scale, $-\log(-\log(F))$, which transforms the tail into a nearly straight line. The 50-year flapwise moment is about 17.8 MNm. The peculiar shape of the extreme load distribution is owed to the fact that it is a mixture of multiple distributions originating from a range of mean wind speeds and different control regions. An example is plotted in Fig. 6, where the data is restricted to $\bar{U} \approx 19$ m/s.

Clearly, the importance sampling method provides a good approximation for the tail of the distribution, given the sample sizes. It also clearly shows the working principle in a qualitative sense. By focusing the computational effort on the extreme loads, it is possible to resolve the high return levels with a small sample size at the cost of having a large error for the lower quantiles. Furthermore, with importance sampling, the tail of the distribution already has its basic shape with a small sample size. This means we do not have to extrapolate that far—or not even at all—which eventually leads to less uncertainty in the final load prediction.

To illustrate this in more detail, subsets were drawn from the existing data sets, each containing one day of simulated time (excluding the start-up period). For every subset, a generalized extreme value distribution was fitted to the data above the 70th percentile ($F > 0.7$) to predict the higher return levels. Fig. 7 then shows the distribution of the 50-year load predictions. Clearly, both methods suffer from a negative bias, owing to the slightly downward curve of the tail of the extreme load distribution. However, the importance sampling method yields a lower uncertainty. Compared to the crude Monte Carlo method, the 90% confidence interval is reduced from [14.8, 19.7] MNm ([–17.1%, +10.7%]) to [16.6, 19.6] MNm ([–7.0%, +10.1%]). In addition, the bias in the predictions decreased from –9.3% to –2.0%.

The best comparison between the two methods is to repeat this for a range of sample sizes as shown in Fig. 8. It shows that the importance sampling method is indeed superior to the crude Monte Carlo simulation over all sample sizes. The crude

Monte Carlo simulation does not produce any usable predictions for a simulated time less than a couple of hours, where less than a handful of data points are actually available for fitting. This where the importance sampling method truly excels over the crude Monte Carlo, since 10 gusts can be evaluated in a single ten-minute period (although, based on these results, we can strongly advise against such small sample sizes). A single measure that compares the two methods directly is the standard error with respect to the "true" 50-year load:

$$\sigma = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(\frac{\widehat{M}_{\mathrm{flap},50,i} - M_{\mathrm{flap},50}}{M_{\mathrm{flap},50}}\right)^2}. \qquad (5)$$

For about 1 month of simulated time, the importance sampling method results in a standard error of 1.4%, as opposed to 6.2% for the crude Monte Carlo simulation.

IV. Discussion

This numerical exercise has led to some useful insights that can be used to further develop the importance sampling method. As is always the case with importance sampling, the quality of the estimate depends on the shape of the sampling distribution. Finding the right sampling distribution can be an iterative process and the effort it takes is not included in the comparison with the crude Monte Carlo simulation. Still, exploring the sample space, as done in Fig. 2, can be rewarding on its own.

It is also difficult to make a really fair comparison between the two methods. A generalized extreme value distribution produced very good fits to the tails of the extreme load distributions. However, it caused a considerable bias that would have been less if we simply used a straight line (a choice we could only make with prior knowledge of the full curve). Especially small sample sizes sometimes resulted in very high loads that, in real life, would have been discarded by a designer. Furthermore, we consistently fitted to the data above the 70th percentile, although slowly shifting this to the 90th or 99th percentile for larger sample sizes could maybe have produced better results.

In retrospect, the reference data set was definitely not ideal for this exercise. The turbulence grid was quite coarse, which meant that there was less room to experiment with large volumetric gusts and gusts targeted at certain parts of the rotor. Moreover, the IEC Kaimal spectrum did not produce the nice coherent structures that would have been obtained from the Mann model. The fact that the gust probability had to be derived by brute force was also far from ideal, since it lead to its own extrapolation problem with uncertainties. Deriving this distribution analytically is clearly far more effective.

V. Conclusions

Preliminary results have shown that predicting extreme loads with importance sampling has several advantages. First and foremost, it has the potential to greatly reduce uncertainty since the computational resources can be efficiently spent on the cases most relevant to the 50-year load. Secondly, the tail of the extreme load distribution already has its basic shape with a small sample size, which makes fitting much easier. Moreover, it removes a large part of the bias that crude Monte Carlo methods can suffer from.

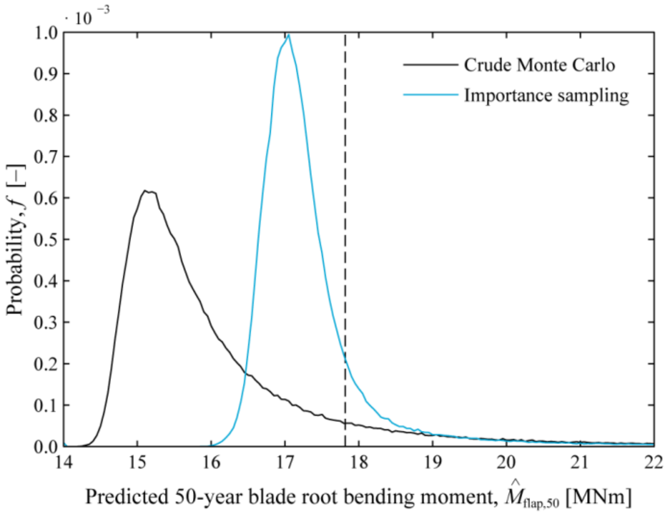

Fig. 7 The uncertainty in the predicted 50-year blade root flapwise bending moments based on 1 day of simulated time. Indicated with a dashed line is the "true" 50-year load based on 96 years of data (17.8 MNm).

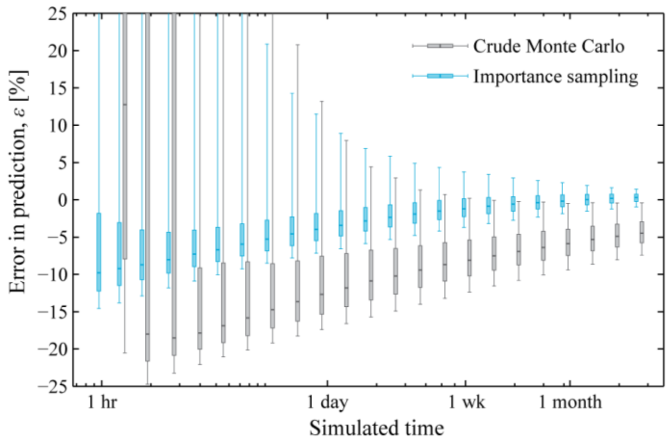

Fig. 8 Box plots showing the error in the prediction of the 50-year blade root flapwise bending moment for both methods. The thick line in the box represents the median, the edges of the box are located at the first and third quartile, and the ends of the whiskers mark the 5th and 95th percentiles.

References

[1] Barone, M., J. Paquette, B. Resor, and L. Manuel (2012). "Decades of wind turbine load simulation". *50th AIAA Aerospace Sciences Meeting*, Nashville, TN, United States. DOI: 10.2514/6.2012-1288.

[2] Bierbooms, W.A.A.M. (2009). "Constrained stochastic simulation of wind gusts for wind turbine design". PhD dissertation, Delft University of Technology.

[3] Bos, R., W.A.A.M. Bierbooms, and G.J.W. van Bussel (2014). 3D stochastic gusts as an alternative to the Mexican hat wavelet. 10th PhD Seminar on Wind Energy in Europe, 28–31 October, Orleans, France.

[4] Bos, R., W.A.A.M. Bierbooms, and G.J.W. van Bussel (2015). "The spatial structure of severe wind gusts". *Manuscript submitted for publication*.

[5] Barone, M., J. Paquette, B. Resor, L. Manuel, and H. Nguyen (2012). "Simulating the entire life of an offshore wind turbine". *EWEA Annual Event*, 16–19 April, Copenhagen, Denmark.

[6] IEC TC88-MT1 (2005). "IEC 61400-1 Ed.3. Wind turbines, part 1: Design requirements". International Electrotechnical Commission, Geneva, Switzerland.

[7] Jonkman, J, S. Butterfield, W. Musial, and G. Scott (2005). "Definition of a 5-MW reference wind turbine for offshore system development", NREL/TP-500-38060. National Renewable Energy Laboratory, Golden, CO., USA.