



Applications of constrained stochastic simulation

Agreement n.: 308974
Duration November 2012 – October 2017
Co-ordinator: DTU Wind

Support by:



PROPRIETARY RIGHTS STATEMENT

This document contains information, which is proprietary to the "INNWIND.EU" Consortium. Neither this document nor the information contained herein shall be used, duplicated or communicated by any means to any third party, in whole or in parts, except with prior written consent of the "INNWIND.EU" consortium.

Document information

Document Name:	XXXXXXXXXXXXXXXXXXXXXX
Document Number:	Deliverable D ...
Author:	XXXXXXXXXXXXXXXXXX
Document Type	Choose an item.
Dissemination level	Choose an item.
Review:	XXXXXXXXXXXXXXXXXX
Date:	XXXXXXXXXXXXXXXXXXXXXX
WP:	XXXXXXXXXXXXXXXXXXXXXX
Task:	XXXXXXXXXXXXXXXXXXXXXX
Approval:	Choose an item.

XXXXXXXXXXXXXXXXXXXXXX

Deliverable D ...

XXXXXXXXXXXXXXXXXX

Choose an item.

Choose an item.

XXXXXXXXXXXXXXXXXX

XXXXXXXXXXXXXXXXXXXXXX

XXXXXXXXXXXXXXXXXXXXXX

XXXXXXXXXXXXXXXXXXXXXX

Choose an item.

1 TABLE OF CONTENTS

1	TABLE OF CONTENTS	5
2	APPLICATIONS OF CONSTRAINED STOCHASTIC SIMULATION	6
2.1	Introduction	6
2.2	Response to extreme gusts.....	6
2.2.1	Set-up.....	6
2.2.2	Response in the yz-plane.....	6
2.2.3	Response in the time domain	7
2.3	Lidar-assisted control	11
2.3.1	Set-up.....	11
2.3.2	Velocity field reconstruction	11
2.3.3	A gust measured in the field.....	13
2.3.4	Construction of a control input signal	14
2.4	DLC 1.1 extreme load predictions	15
2.4.1	Case study: the NREL 5 MW.....	15
2.4.2	An importance sampling method	17
2.5	References	20
3	CHAPTER 1	21
3.1	First paragraph.....	21
3.1.1	Second paragraph.....	21
4	CHAPTER 2	22
5	CHAPTER 3	23
5.1	First Paragraph.....	23
5.1.1	xxxx.....	23

2 APPLICATIONS OF CONSTRAINED STOCHASTIC SIMULATION

2.1 Introduction

Constrained stochastic simulation (D1.12) is a method that allows the user to generate extreme events in 3D turbulent wind fields. In addition, the probability associated with local maxima and minima can be derived analytically. This means that load cases containing the 50-year gust can be set up directly, without having to generate 50 years of turbulence.

Having full control over the wind field offers great potential. First, it makes it possible to study the interaction between the rotor and extreme gusts (which would otherwise naturally occur in very long time series). Second, it can be used to reconstruct the velocity field around lidar measurements. And, finally, they open the door for very efficient extreme load predictions when coupled to importance sampling methods. These applications will be discussed in detail in the following sections.

What is presented here is a brief overview of the work conducted in the scope of D1.43. The full extent of TU Delft's contribution to D1.12 and D1.43 is compiled in a PhD thesis:

Bos, R. (2017). *Extreme gusts and their role in wind turbine design*. PhD thesis. Delft University of Technology.

For completeness, in Appendix A of this report Chapter 4 of the PhD thesis is included: Gust loads on rotor blades; Appendix B is Chapter 5: Predicting extreme gust loads.

2.2 Response to extreme gusts

Compared to, for example, the IEC extreme operating gust, the gusts generated by constrained stochastic simulation are bounded in space and can be positioned anywhere in the wind field. To assess the impact this has on a turbine—most notably the blade root flapwise moments—a large number of these gusts are generated at various locations on the rotor disk.

2.2.1 Set-up

As an example, we consider the DTU 10 MW with the baseline controller in Bladed v4.4. Based on the steady thrust and pitch curve, we expect that the rated wind speed of $\bar{U} = 11.4$ m/s is the point at which the turbine is most susceptible to high gust loads. A severe gust is set up by taking a spheroidal volume with a length of $2\bar{U}$ and a diameter of 25 m. Based on the theory of D1.12, this corresponds to a 50-year amplitude of 11.0 m/s.

A 64×64 , 4-Hz grid is set up with a width and height of 268.8 m ($= 8L$). Each simulation is run for 2 minutes, of which the first minute is discarded to avoid the start-up transients. The gust arrives around the 24-s mark, but varies within a period of $(2\pi/\Omega)/3$ to make sure that every possible blade position is accounted for. The gusts are positioned uniformly over the frontal plane at 32×32 positions, each with 9 seeds, resulting in a total of 9,216 simulations.

2.2.2 Response in the yz-plane

Figure 1 shows the tower overturning and blade flapwise bending moments as a function of the gust landing on a range of y - and z -positions. In both cases, the upper half of the rotor disk is the most vulnerable part, which is due to the wind shear profile. When added to the gust velocity, it causes high apparent wind speeds and higher angles of attack near the zero azimuth position. For the tower overturning moment, the gusts landing higher up have an even bigger impact due to the longer arm with respect to the tower base.

In addition, there is also some asymmetry in the lateral plane. We investigated this by feeding the mean gust shape to a simplified, fixed speed, fixed pitch model of the DTU 10 MW rotor. This mean gust shape is retrieved from the expected value of Equation (2.5) of the D1.12 report; that is,

$$\mathbf{n}_c = \mathbf{A}^*(\mathbf{A}\mathbf{A}^*)^{-1}\mathbf{b}. \quad (2.1)$$

Using the rotationally sampled velocity signals, we found that the asymmetry is caused by the correlation tensor prescribed by the Mann model. Positive streamwise amplitudes are often accompanied by negative (downward) vertical velocities. Therefore, a blade hit by a gust during its downstroke will also experience a decrease in the azimuthal velocity component. This causes

stronger angle of attack changes with higher (dynamic) lift coefficients. Despite the fact that the total velocity component is slightly reduced, the total lift force on the blade is higher. Of course, the opposite happens during an upstroke. The result is that the right side of the rotor (when viewed from upwind) is slightly more susceptible to high gust loads.

Still, there are a lot of positions—think of the blade root—where a severe gust does not trigger any significant load whatsoever. This has an effect on the 50-year load level, since the probability of a gust landing on a certain position has to be included in the load calculation. Therefore, the expected flapwise bending moment associated with this 50-year event (i.e., the risk) is not very high (only 41.0 MN·m). The actual 50-year load level cannot be obtained directly from a 50-year gust, however. It is an extensive process where a designer has to take into account the contributions of weaker and stronger gusts, landing at every imaginable position. This is discussed in Section 2.4.

2.2.3 Response in the time domain

What happens in the time domain is plotted in Figure 2 and Figure 3. The maximum bending moments lag behind the maximum amplitude by roughly 5 seconds, although at that time the longitudinal wind speed is still well above the mean. In most cases (at least 95%), the baseline controller handles the gust well (see Figure 2b and c). However, there are cases where the gust does trigger high bending moments. One of those is depicted in Figure 3. During that specific event, the rotor was running below the rated speed. This meant that there was still some room for the rotor to speed up before the pitch system had to intervene. The controller indeed managed to limit the torque overshoot, but not before the turbine experienced high maximum bending moments.

Compared to gusts with a fully uniform inflow over the yz -plane (e.g., the IEC extreme operating gust), there are some important differences. Firstly, a local velocity peak also means that the gust loads do not follow the time signature of the gust, as would be the case with the IEC extreme operating gust. Instead, blade loads and tower loads respectively follow 1P and 3P frequencies.

It is difficult to identify gusts and limit loads based on the high-speed shaft (HSS) torque or other rotor-effective input signals. With blades the size of the DTU 10 MW, gusts are local velocity peaks that often only affect one blade at a time. Therefore, a high bending moment acting on a

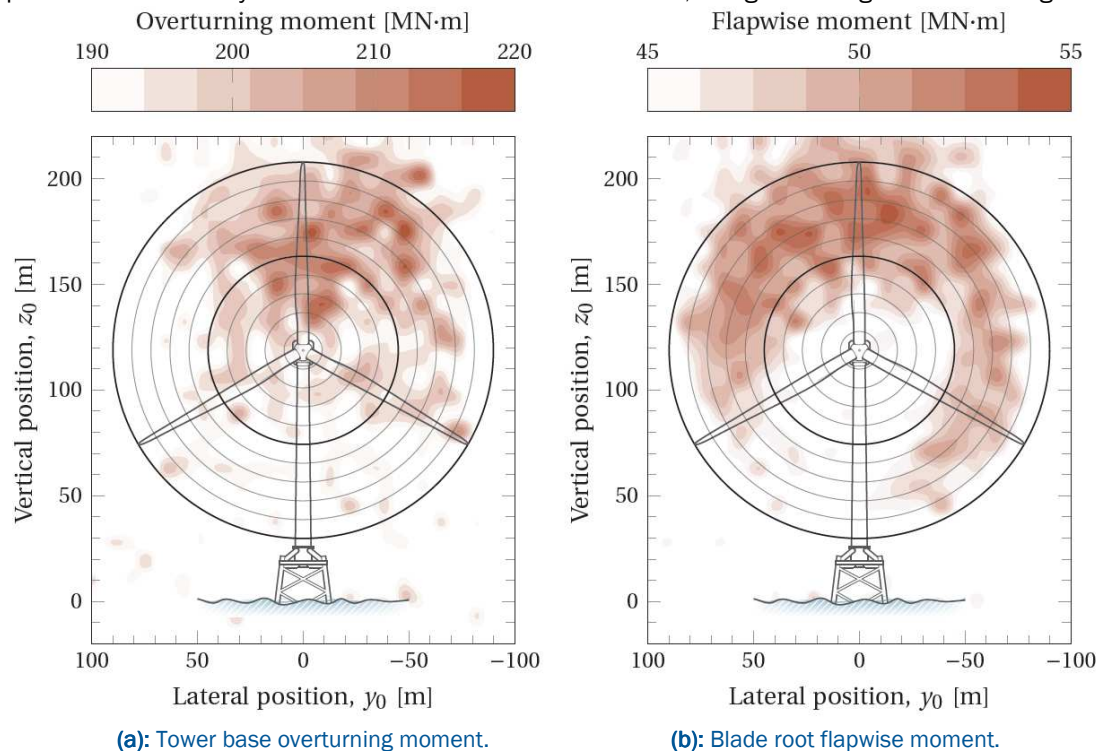




Figure 1: Bending loads (indicated by color, see color bar at the top), depending on the lateral and vertical gust position (indicated by the y - z position in the rotor plane).

single blade root is not necessarily accompanied by high torque. For the same reason, high blade root moments do not always mean high tower moments.


Because the angle of attack changes are very local, collective pitch is a very crude measure and can lead to adverse results in the case of negative gust amplitudes. In the case of the NREL 5 MW machine with the baseline controller (see Figure 9), this turned out to be solely responsible for the 50-year load. A better solution is to rely on individual pitch, distributed control, and/or lidar-assisted control.



(a): Longitudinal wind speed with the gust centers marked with the shaded area. Here, the gust is centered on the hub at 24 s.



(b): Medians, confidence intervals and maxima of the tower base overturning moments.



(c): Medians, confidence intervals and maxima of the blade root flapwise bending moments.

Figure 2: Time series of all 9,216 simulated load cases.

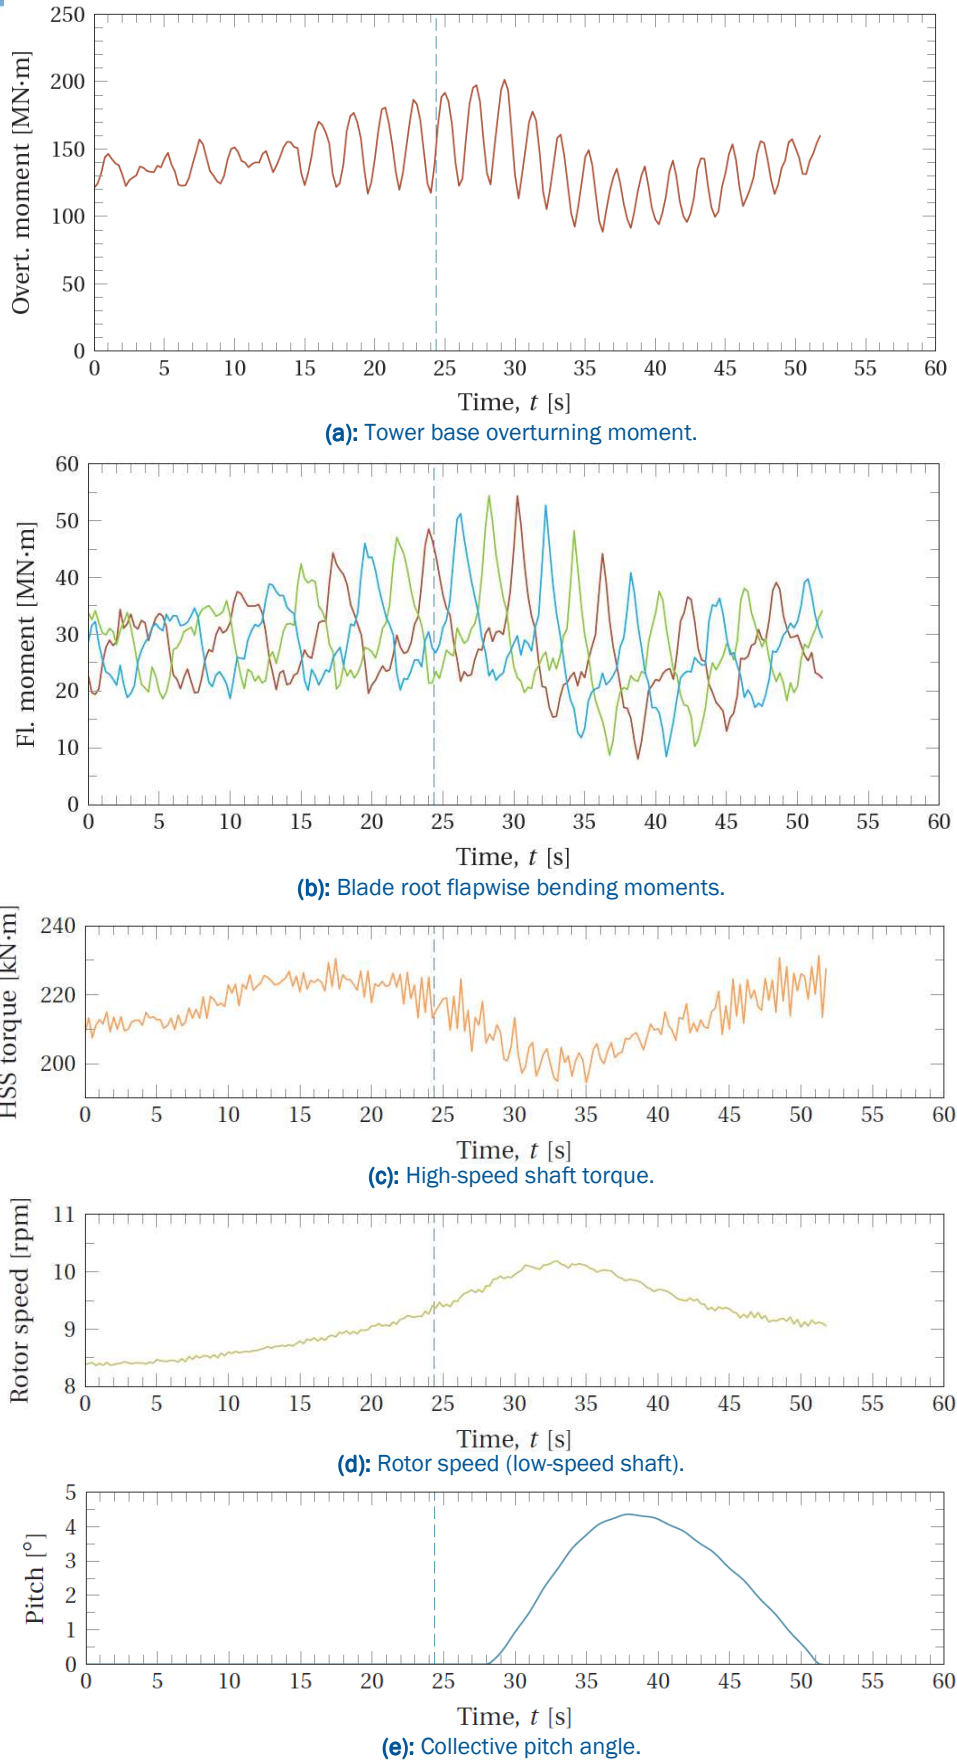


Figure 3: Time series belonging to the maximum bending moment (55.4 MN·m, with the extreme gust positioned at $t_0 = 24.5$ s, $y_0 = -13$ m, $z_0 = 184$ m).




Figure 4: Sketch of the pulsed lidar system mounted on the nacelle of a Darwind XD115-5MW wind turbine.

2.3 Lidar-assisted control

Another interesting application is the reconstruction of lidar data. On the assumption that turbulence is a stationary, homogeneous, and Gaussian process, the statistics of any 3D spectral model can be used to give the best possible estimate for the velocity field. This was demonstrated using data from the LAWINE project (concluded in October 2016). More details of this method can be found in the paper (Bos *et al.*, 2016).

2.3.1 Set-up

The set-up consisted of an Avent-Lidar 5-beam pulsed lidar prototype that was mounted on the nacelle of a Darwind XD115-5MW wind turbine on the test site of ECN (see Figure 4). It provides data of ten range gates simultaneously (50–185 m upwind), while cycling between the five beam positions (0–4) every 0.25 s.

The lidar does not provide a perfect point measurement, but more of a weighted average. This is modeled by a Gaussian pulse shape (Frehlich *et al.*, 2006):

$$w(r) = \frac{1}{2\Delta R} \left[\operatorname{erf} \left(\frac{R - r + \frac{1}{2}\Delta R}{\Delta Z} \right) - \operatorname{erf} \left(\frac{R - r - \frac{1}{2}\Delta R}{\Delta Z} \right) \right] \quad (2.2)$$

where R is the distance to the range gate, ΔR the distance between the range gates, ΔZ the e^{-1} radius of the pulse, and $\operatorname{erf}(x)$ is the error function. The e^{-1} radius is derived from the full-width-at-half-maximum pulse width (FWHM):

$$\Delta Z = \frac{\text{FWHM}}{2\sqrt{\ln 2}}, \quad (2.3)$$

with $\text{FWHM} = 30$ m for this set-up. Under the assumption that the streamwise velocity component is the dominant one (i.e., $u \gg v, w$), the line-of-sight velocity at a given point j is:

$$u_{\text{LOS},j} = \cos \varphi_j \int u(x_j + r, y_j, z_j) w(r) dr, \quad (2.4)$$

where φ is the beam angle.

2.3.2 Velocity field reconstruction

The 3D velocity field that surrounds the collection of n measurements is a conditional field:

$$\hat{\mathbf{u}}(\mathbf{x}) = \left\{ \mathbf{u}(\mathbf{x}) \mid \int u(x_1 + r, y_1, z_1) w(r) dr = \frac{u_{\text{LOS},1}}{\cos \varphi_1}, \right. \\ \left. \dots, \int u(x_n + r, y_n, z_n) w(r) dr = \frac{u_{\text{LOS},n}}{\cos \varphi_n} \right\}. \quad (2.5)$$

Using the theory of D1.12, it is possible to derive the set of velocity fields that match the set of measurements, while adhering to the statistics of the spectral tensor, $\Phi(\boldsymbol{\kappa})$. For this, we assume that $\mathbf{u}(\mathbf{x})$ is stationary, homogeneous, and Gaussian and can be constructed by a Fourier series:

$$\mathbf{u}(\mathbf{x}) = \bar{\mathbf{u}}(\mathbf{x}) + \sum_{\boldsymbol{\kappa}} \mathbf{C}(\boldsymbol{\kappa}) \mathbf{n}(\boldsymbol{\kappa}) e^{i\boldsymbol{\kappa} \cdot \mathbf{x}}, \quad (2.6)$$

where $\mathbf{u} = [u, v, w]^T$ is a velocity vector, $\mathbf{x} = [x, y, z]^T$ a position vector, $\boldsymbol{\kappa} = [\kappa_x, \kappa_y, \kappa_z]^T$ the wave number vector, $\mathbf{C}(\boldsymbol{\kappa})$ a correlation tensor (obtained by a Cholesky decomposition of the spectral tensor; e.g., see (B.13) of IEC 61400-1, 2nd edn., Appendix B), and $\mathbf{n}(\boldsymbol{\kappa}) \sim CN(0, \mathbf{I}_3)$ a vector of complex-normal distributed coefficients. Moreover, $\bar{\mathbf{u}}(\mathbf{x})$ denotes the (time-invariant) mean wind speed component. This can be written as a matrix multiplication according to:

$$\mathbf{u}(\mathbf{x}) - \bar{\mathbf{u}}(\mathbf{x}) = \boldsymbol{\Psi} \mathbf{n}, \quad (2.7)$$

where

$$\boldsymbol{\Psi} = [\dots, \mathbf{C}(\boldsymbol{\kappa}_{j-1}) e^{i\boldsymbol{\kappa}_{j-1} \cdot \mathbf{x}}, \mathbf{C}(\boldsymbol{\kappa}_j) e^{i\boldsymbol{\kappa}_j \cdot \mathbf{x}}, \mathbf{C}(\boldsymbol{\kappa}_{j+1}) e^{i\boldsymbol{\kappa}_{j+1} \cdot \mathbf{x}}, \dots], \quad (2.8)$$

is a Fourier transform matrix and where

$$\mathbf{n} = \begin{bmatrix} \vdots \\ \mathbf{n}(\boldsymbol{\kappa}_{j-1}) \\ \mathbf{n}(\boldsymbol{\kappa}_j) \\ \mathbf{n}(\boldsymbol{\kappa}_{j+1}) \\ \vdots \end{bmatrix}, \quad (2.9)$$

is a one-dimensional white noise vector. From (2.8), the streamwise velocity component is easily obtained by taking the first row; i.e.,

$$u(\mathbf{x}) - \bar{u}(\mathbf{x}) = \boldsymbol{\Psi}_u \mathbf{n}, \quad (2.10)$$

where

$$\boldsymbol{\Psi}_u = [\dots, [C_{uu}(\boldsymbol{\kappa}_j), C_{uv}(\boldsymbol{\kappa}_j), C_{uw}(\boldsymbol{\kappa}_j)] e^{i\boldsymbol{\kappa}_j \cdot \mathbf{x}}, \dots]. \quad (2.11)$$

What now follows is the same approach used in Section 2.4 of the D1.12 report. When the complete set of lidar measurements, $1, \dots, n$, are stored in a one-dimensional vector, \mathbf{b} :

$$\mathbf{b} = \begin{bmatrix} u_{\text{LOS},1} / \cos \varphi_1 - \bar{u}(\mathbf{x}_1) \\ u_{\text{LOS},2} / \cos \varphi_2 - \bar{u}(\mathbf{x}_2) \\ \vdots \\ u_{\text{LOS},n} / \cos \varphi_n - \bar{u}(\mathbf{x}_n) \end{bmatrix}, \quad (2.12)$$

Equation (2.5) can be rewritten as a linear system:

$$\hat{\mathbf{u}}(\mathbf{x}) = \boldsymbol{\Psi} \{ \mathbf{n} | \mathbf{Y} \mathbf{n} = \mathbf{b} \}, \quad (2.13)$$

where the matrix \mathbf{Y} follows from stacking the matrix (2.11) and filtering it with the kernel $w(r)$:

$$\mathbf{Y} = \begin{bmatrix} \dots, W(\kappa_j) [C_{uu}(\boldsymbol{\kappa}_j), C_{uv}(\boldsymbol{\kappa}_j), C_{uw}(\boldsymbol{\kappa}_j)] e^{i\boldsymbol{\kappa}_j \cdot \mathbf{x}_1}, \dots \\ \dots, W(\kappa_j) [C_{uu}(\boldsymbol{\kappa}_j), C_{uv}(\boldsymbol{\kappa}_j), C_{uw}(\boldsymbol{\kappa}_j)] e^{i\boldsymbol{\kappa}_j \cdot \mathbf{x}_2}, \dots \\ \vdots \\ \dots, W(\kappa_j) [C_{uu}(\boldsymbol{\kappa}_j), C_{uv}(\boldsymbol{\kappa}_j), C_{uw}(\boldsymbol{\kappa}_j)] e^{i\boldsymbol{\kappa}_j \cdot \mathbf{x}_n}, \dots \end{bmatrix}. \quad (2.14)$$

where $W(\kappa_j)$ is the Fourier transform of $w(r)$.

Now, under the strict assumption of Gaussianity, the statistics of $\hat{\mathbf{u}}(\mathbf{x})$ can be derived analytically:

$$E[\hat{\mathbf{u}}(\mathbf{x})] = \bar{\mathbf{u}}(\mathbf{x}) + \boldsymbol{\Psi} \mathbf{Y}^* (\mathbf{Y} \mathbf{Y}^*)^{-1} \mathbf{b}, \quad (2.15)$$

$$\text{var}[\hat{\mathbf{u}}(\mathbf{x})] = \boldsymbol{\Psi} \boldsymbol{\Psi}^* - \boldsymbol{\Psi} \mathbf{Y}^* (\mathbf{Y} \mathbf{Y}^*)^{-1} \mathbf{Y} \boldsymbol{\Psi}^*, \quad (2.16)$$

where $*$ denotes a conjugate transpose. This means that $E[\hat{\mathbf{u}}(\mathbf{x})]$ is the expected (mean) velocity field surrounding the measurements. Furthermore, $\text{var}[\hat{\mathbf{u}}(\mathbf{x})]$ is the variance. At the measurement points, the variance is zero (i.e., the velocity is always equal to the measurement), while infinitely

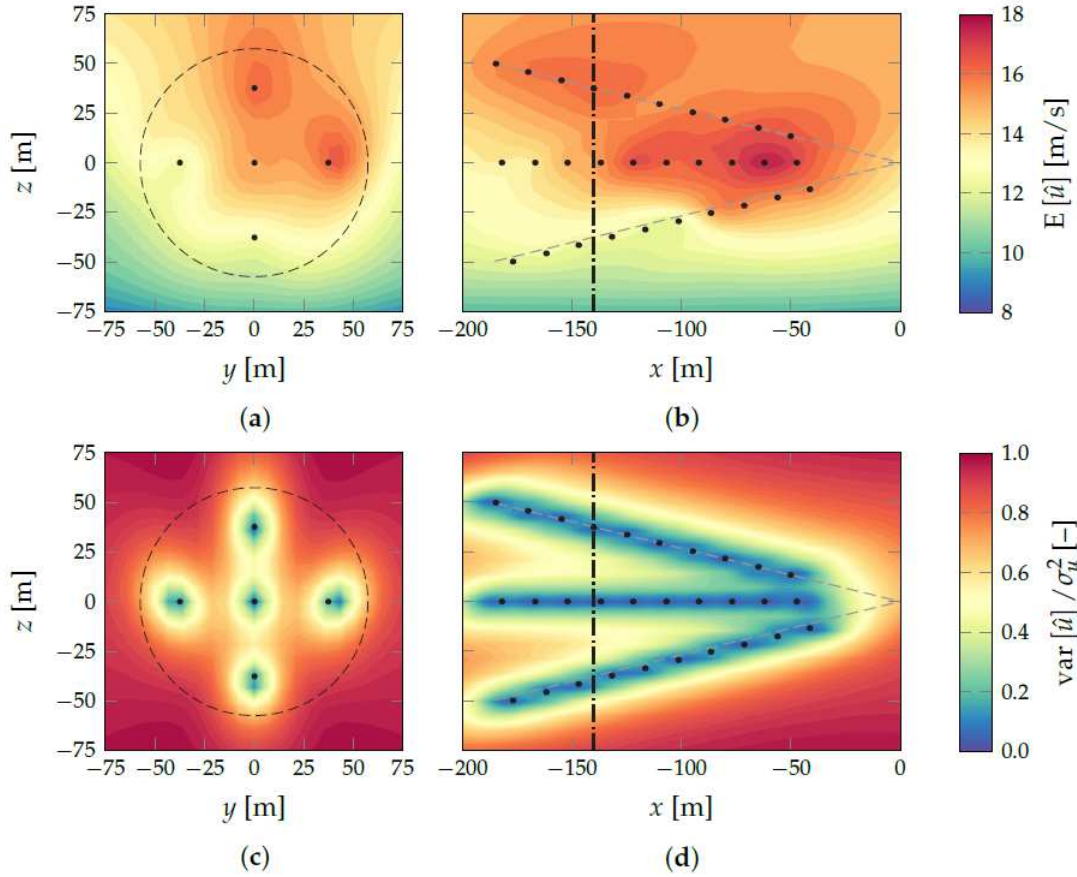


Figure 5: Expected velocity fields in the yz- (a) and xz-planes (b), together with the normalized variance in the yz- (c) and xz-planes (d).

far away, the variance should be equal to the square of the turbulence intensity (i.e., the best guess with only the ten-minute statistics). The process can be best compared to *kriging*—which was originally applied in geostatistics—where the variance is representative of the uncertainty.

2.3.3 A gust measured in the field

This method was applied to a gust that was picked up by the lidar's center beam on 22 December 2013. It was the most severe event in the 1.5-year data set, while still being undisturbed from the surrounding turbines and met masts.

Assuming that the terrain is relatively flat and homogeneous, a neutral wind shear profile was fitted to the ten-minute mean wind speeds measured at the positions 0, 1 and 3 at $r = -140$ m:

$$\bar{u}(z) = \bar{u}(z_{\text{ref}}) \frac{\ln(z/z_0)}{\ln(z/z_{\text{ref}})}, \quad (2.17)$$

which resulted in a roughness length of $z_0 = 0.55$ m, matching the site fairly well (de Jong *et al.*, 1990). In addition, Mann's spectral tensor was set up with $L \approx 25$ m (estimated from the low-frequency part of the lidar spectrum), $\Gamma = 3$ (for neutral conditions) (Sathe *et al.*, 2013), and $\alpha \epsilon^{2/3} \approx 0.16$ m^{4/3}/s². The latter was found by matching the longitudinal variance of the center beam measurements to the filtered spectral tensor:

$$\int W^2(\kappa_x) \Phi_{uu}(\boldsymbol{\kappa}) d\boldsymbol{\kappa} \approx \sigma_{u, \text{LOS}}^2, \quad (2.18)$$

Figure 5 shows two cross-sections of the fields $E[\hat{\mathbf{u}}(\mathbf{x})]$ and $\text{var}[\hat{\mathbf{u}}(\mathbf{x})]$. The low-frequency parts of the gust, which hold most of the momentum content, are captured fairly well. The variance is the lowest along the directions of the beams, as should be expected. However, it never reaches complete zero due to the range weighting, so part of this variance has to account for the high frequencies that are not captured by the lidar.

2.3.4 Construction of a control input signal

The expected velocity field, together with its variance, can be used to construct an input signal for a gust controller. It relies on the notion that, under Taylor's frozen turbulence hypothesis, it holds that

$$\begin{aligned}\rho \bar{u} \frac{\partial u}{\partial x} &= -\frac{\partial p}{\partial x}, \\ \frac{\partial}{\partial x}(\rho \bar{u} u + p) &= 0, \\ \rho \bar{u} u + p &= \text{const.}\end{aligned}\tag{2.19}$$

Here, ρ is the air density and p is the static atmospheric pressure. The along-wind force, pushing on the yz -plane at a distance r upwind, is found by integrating the dynamic pressure term, $\rho \bar{u} u$, over a surface, A :

$$F(t; r) = \rho \iint_A \bar{u}(r, y, z) u(r, y, z, t) dy dz.\tag{2.20}$$

Because $\bar{u}(\mathbf{x})$ is constant and $u(\mathbf{x}, t)$ is a Gaussian field, it must follow that $F(t; r)$ is also Gaussian. This is convenient, since it allows us to also derive its statistics, based on the statistics of a reconstructed velocity field, $\hat{u}(\mathbf{x}, t)$:

$$E[\hat{F}(t; r)] = \rho \iint_A \bar{u}(r, y, z) E[\hat{u}(r, y, z, t)] dy dz,\tag{2.21}$$

$$\text{var}[\hat{F}(t; r)] = \rho^2 A \iint_A \bar{u}^2(r, y, z) \text{var}[\hat{u}(r, y, z, t)] dy dz.\tag{2.22}$$

This along-wind force is related to the loads on the structure and can be used for gust control. For example, with 97.7% certainty, the along-wind force will stay under the level $\hat{F} + 2\sigma_{\hat{F}}$ (see Figure 6). The uncertainty is a function of the turbulence coherence and the amount of available measurement points, which will vary due to the technical availability of the lidar and the blades passing in front of the eye. Therefore, a signal $\hat{F} + n\sigma_{\hat{F}}$ is a rather objective way of judging the lidar input.

The method was cross-checked with a similar input signal: the rotor-effective wind speed, U_{eff} , which is a pseudo-signal that is directly derived from the measured rotor torque:

$$Q = \frac{1}{16} \rho C_Q(\theta, \lambda) \pi D^3 U_{\text{eff}}^2,\tag{2.23}$$

where Q is the shaft torque, C_Q the torque coefficient (a function of the pitch angle, θ , and tip speed ratio, λ), and D the rotor diameter. The rotor-effective wind speed has been proven to be a good input signal for feed-forward control and has been successfully implemented to improve turbine behavior under gust loading (van der Hooft *et al.*, 2003, 2004). Figure 7 shows a direct comparison of the two signals by matching their mains. The rotor-effective wind speed was downsampled to 1 Hz (was 64 Hz) to reduce the noise and focus on the more important low-frequency fluctuations. The 95% range follows from the area enclosed by $\hat{F} \pm 2\sigma_{\hat{F}}$. The two signals seem to agree well in a qualitative sense, with the most important up- and downward trends around $t = 0$ clearly recognizable.




Figure 6: Expected along-wind force belonging to the gust event of Figure 5 at $r = -140$ m, together with the uncertainty levels.




Figure 7: Expected along-wind force belonging to the gust event of Figure 5 at $r = -140$ m, compared to the rotor-effective wind speed by ECN with a 10.5-s delay to account for advection towards the rotor disk (downsampled to 1 Hz for clarity).

2.4 DLC 1.1 extreme load predictions

Probably the most powerful application of constrained stochastic simulation is in the prediction of extreme loads. Extreme loads, most notably the 50-year loads, are difficult to obtain because certain dynamic behavior or controller behavior might only show up in very large sample sizes. This often makes it hard to model the extreme load behavior in cases when the computational budget is limited; for example, in early design phases.

2.4.1 Case study: the NREL 5 MW

As a case study, we considered the NREL 5 MW. This was a convenient subject, since a large data set was available with 96 years' worth of extremes (Barone *et al.*, 2013).¹ At the same time, it provided a nice insight in the controller behavior in extreme cases. The data set was generated

¹ Downloadable from http://energy.sandia.gov/?page_id=13173.

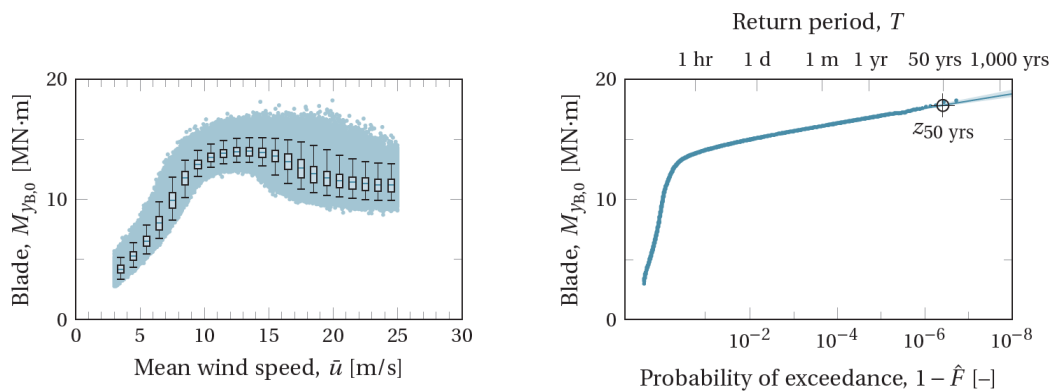
using the onshore version of the NREL 5 MW with the baseline controller, modeled in FAST v7. Wind fields were generated by TurbSim on a 137×137 m, 20×20 grid, with a temporal frequency of 20 Hz, and following an IEC class 1B normal turbulence model. Each simulation was run for 11 minutes, with the first minute discarded to get rid of the start-up transients.

The data set was generated by a crude Monte Carlo method. This means that the mean wind speeds were sampled from the parent Rayleigh distribution and that the extreme load distribution follows naturally from sorting and assigning the appropriate plotting position to the i th load:

$$\hat{F}(M_{y,i}) = \frac{i}{N + 1}, \quad (2.24)$$

where N is the sample size. The regions below the cut-in and above cut-out wind speeds can be accounted for by either completing the data set with zero loads or by correcting the load distribution according to

$$\hat{F}(M_{y,i}) = 1 - \left(1 - \frac{i}{N + 1}\right) \int_{\text{cut-in}}^{\text{cut-out}} f(\bar{U}) d\bar{U}. \quad (2.25)$$



(a): Scatter plot of the entire data set. The edges of the box plots indicate the 25th and 75th percentiles, the whiskers indicate the 2.5th and 97.5th percentiles, and the bar shows the medians.

(b): Return level plot, yielding a 50-year extreme of 17.8 MN·m. Longer return periods are obtained by fitting a generalized extreme value distribution. The shaded area marks the 95% confidence interval, which is estimated by resampling the data.

Figure 8: Extreme blade root flapwise moments of the NREL 5 MW reference turbine ($N = 5 \cdot 10^6$).

Figure 8 shows the extreme blade root flapwise bending moments. Judging from both the scatter plot (a) and the return level plot (b), the extreme loads seem to originate from relatively high wind speeds in the pitch control regime. It turned out that this was due to a particular weakness in the baseline controller. Strongly negative gust amplitudes would sometimes reduce the wind speed close to the rated wind speed, which causes the turbine to pitch back to zero. When the wind speed recovers, the machine is operating at full thrust in an 18–19-m/s mean wind speed, leading to the high extreme loads (e.g., see Figure 9). When the data set increases in size, the extreme cases appear in even higher wind speeds, since increasingly larger amplitudes are encountered that are able to bring the wind speed back to rated. This also means that, for small sample sizes, the extremes might seem to lie close to the rated wind speed, which is also what is shown by the error bars in Figure 8a.

2.4.2 An importance sampling method

Not only is it very computationally intensive to establish a 50-year load with sufficient accuracy, it is also very difficult to quantify whether an updated control strategy is able to reduce the 50-year load. A possible solution is to rely on an importance sampling method that uses constrained gusts. The extreme events that cause the behavior shown in Figure 9 can be sampled directly from a distribution $g(\boldsymbol{\theta})$. An extreme load distribution then follows from

$$\hat{F}(M_y) = \frac{\sum_{i=1}^N \mathbf{1}(M_{y,i} \leq M_y) \frac{f(\boldsymbol{\theta}_i)}{g(\boldsymbol{\theta}_i)}}{\sum_{i=1}^N \frac{f(\boldsymbol{\theta}_i)}{g(\boldsymbol{\theta}_i)}}, \quad (2.26)$$

where the probabilities are weighted by a likelihood ratio, $f(\boldsymbol{\theta})/g(\boldsymbol{\theta})$. Here, $\boldsymbol{\theta}$ is a vector in a parameter space, $f(\boldsymbol{\theta})$ is the parent distribution, $g(\boldsymbol{\theta})$ is a sampling distribution, and

$$\mathbf{1}(x \in S) = \begin{cases} 1, & \text{if } x \in S, \\ 0, & \text{if } x \notin S, \end{cases} \quad (2.27)$$

is the indicator function.

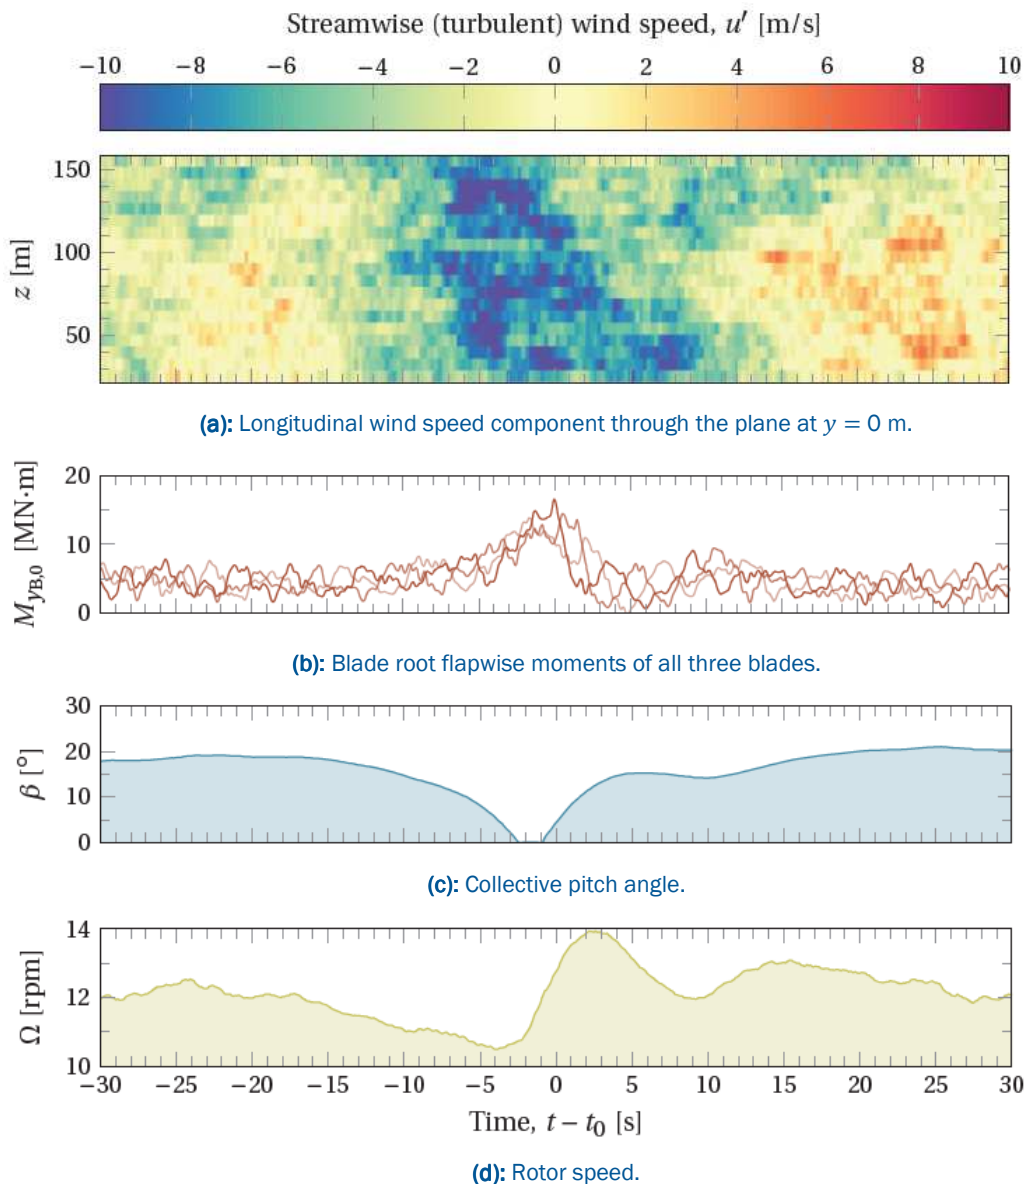


Figure 9: Time series belonging to an extreme flapwise bending moment at $\bar{U} = 19$ m/s.

Importance sampling can be a very attractive solution when a designer has complete control over the wind field. Instead of running large sample sizes of ten-minute wind fields, the same results can be obtained by only evaluating the response to a small set of extreme gusts, which can be stored in time series of 1–2 minutes.

The critical part of an importance sampling method is always to define a proper sampling distribution. We assume that the extremes loads are dependent on 2 parameters: the mean wind speed, \bar{U} , and the gust amplitude, A . The gust position is left unconstrained (i.e., uniformly distributed over the yz -plane). Then, an initial survey of the parameter space, shown in Figure 10, explains the scatter of Figure 8a. It confirms the suspicion that combinations of high wind speeds and large negative gust amplitudes are responsible for the extreme load behavior.




Figure 10: Extreme blade root flapwise moments of the NREL 5 MW reference turbine triggered by point gust amplitudes ($N = 10^5$).

The probability associated with such gusts are estimated through the *Euler characteristic heuristic*:

$$f(A) \approx -\frac{d}{dA} E[\varphi(Z_A)], \quad (2.28)$$

where the expected Euler characteristic, $E[\varphi(Z_A)]$, is given by Equation (3.13) of the D1.12 report. Based on Figure 10, two different gust settings are used: (1) a point gust (i.e., exciting a single grid point) and (2) a spheroidal gust with a time scale of 2 s and a lateral diameter of 25 m. The 50-year amplitudes for these events lie approximately at the $\pm 8\sigma$ and $\pm 7\sigma$ level, respectively. Two sampling distributions are then set up:

$$g_1(\bar{U}, A) \approx \frac{1}{2\pi \cdot 2 \cdot 0.5} \exp \left[-\frac{(U - 20)^2}{2 \cdot 2^2} - \frac{(A/\sigma + 7)^2}{2 \cdot 0.5^2} \right], \quad (2.29)$$

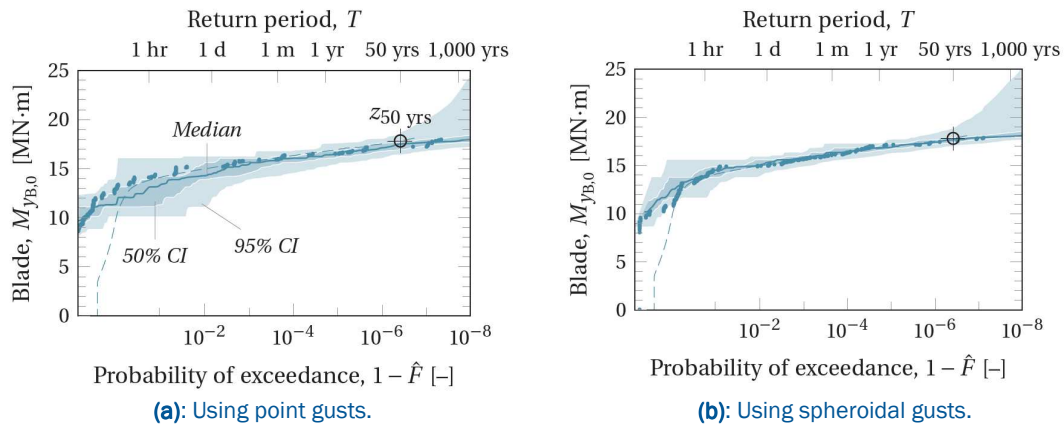


Figure 11: Return level plots obtained by importance sampling ($N = 10^4$). The confidence intervals (CI) were estimated by resampling a larger set of $5 \cdot 10^4$ gust loads. The dashed line indicates the extreme load distribution constructed from the full, 96-year data set. The dot markers belong to an arbitrary sample.

$$g_2(\bar{U}, A) \approx \frac{1}{2\pi \cdot 2 \cdot 0.5} \exp \left[-\frac{(U - 20)^2}{2 \cdot 2^2} - \frac{(A/\sigma + 6)^2}{2 \cdot 0.5^2} \right], \quad (2.30)$$

which are two normal distributions centered on a slightly lower level than the 50-year level. Gusts are then generated from the Kaimal spectrum in two-minute time series (of which the first minute was used as a start-up buffer) and using the same grid settings as in the reference data set. The loads were obtained through FAST v7.²

Figure 11 shows the return level plots obtained by these sampling distributions on the basis of a sample size of $N = 10^4$. The 50-year load can often be interpolated, which eliminates the need of an extrapolation scheme. With both gust types, the 50-year load is obtained with high accuracy, although the spheroidal gusts produce better results. This is because they are larger and contain more momentum. Therefore, they have a bigger impact on the loads and are thus better correlated to the extremes.³ The point gusts produce a 95% confidence interval of [16.4, 18.6] MN·m (-7.9%, +4.5%). The spheroidal gusts [17.0, 18.8] MN·m (-4.5%, +5.6%).

An advantage of working with gusts is that they can be stored in one-minute time series. When ignoring the start-up buffer, a sample size of $N = 10^4$ gusts has the same computational burden of $N = 10^3$ ten-minute wind fields. In order to make a direct comparison, the performance of the importance sampling method was compared to that of a conventional method that relies on sampling ten-minute wind speeds, either directly from the parent Rayleigh distribution or from a uniform distribution bounded by the cut-in and cut-out points. The extrapolation scheme was automated by setting the starting point of the distribution tail as the center of the plot on Gumbel paper; i.e.,

$$-\ln[-\ln(\hat{F})] > \frac{1}{2} \min\{-\ln[-\ln(\hat{F})]\} + \frac{1}{2} \max\{-\ln[-\ln(\hat{F})]\}. \quad (2.31)$$

This produced results that were fairly accurate but, most of all, consistent and objective. Extrapolation was then done by fitting a generalized extreme value distribution to the tail:

$$G(M_y; \mu, \sigma, \xi) = \exp \left[-\left(1 + \xi \frac{M_y - \mu}{\sigma}\right)^{-\frac{1}{\xi}} \right], \quad (2.32)$$

which seemed to match the tail best for large sample sizes.⁴

The performance of the methods is expressed as the median absolute deviation to the true 50-year bending moment of 17.8 MN·m:

$$\varepsilon_{\text{MAD}} = \text{median}(|\hat{M}_y - M_y|). \quad (2.33)$$

² A prior check was run with a large number of ten-minute, unconstrained fields to confirm that they produce the same results as the reference data set.

³ In principle, any gust shape should be able to produce the correct return levels.

⁴ For small sample sizes, a straight line often produced better results.

This is similar to the root-mean-squared error, but is less distorted by extremely bad fits. Figure 12 then shows the comparison of the methods. For the same quality result, the importance sampling method is, in general, about two orders of magnitude faster than sampling ten-minute mean wind speeds. This makes it especially suitable for early design phases when computational resources are scarce.

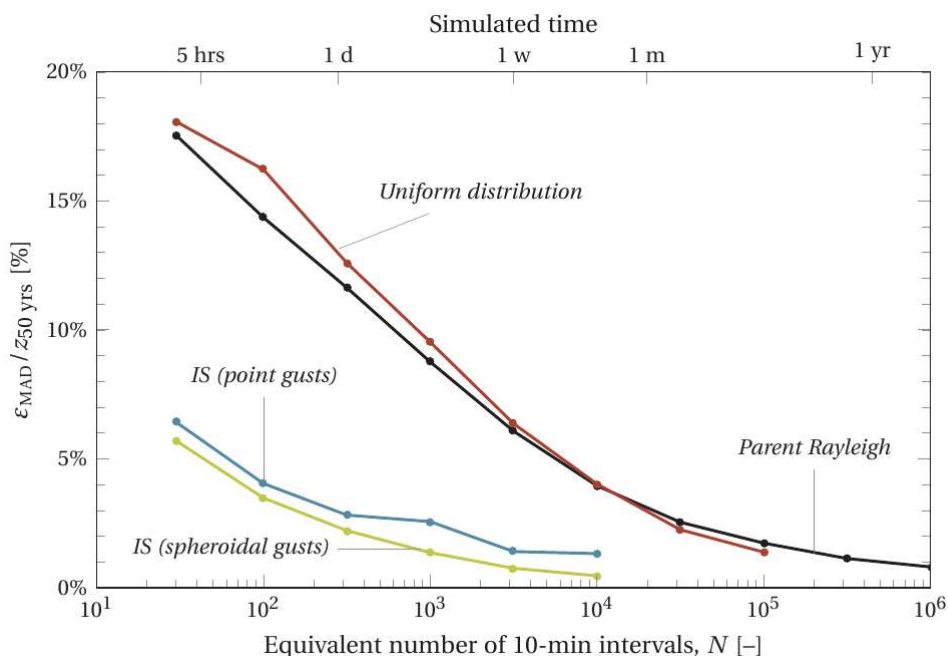


Figure 12: Median-absolute deviation to the 50-year load as obtained by the importance sampling method (IS), compared to sampling ten-minute mean wind speeds from a uniform distribution and from the parent Rayleigh distribution (i.e., crude Monte Carlo).

2.5 References

- Barone, M., Paquette, J., Resor, B., and Manuel, L. (2012). Decades of wind turbine load simulation, 50th AIAA Aerosp. Sci. Meet. Incl. New Horizons Forum Aerosp. Expo., Nashville, TN, United States, 9–12 January, 2012, doi:10.2514/6.2012-1288.
- Bos R, Giyanani A, Bierbooms W. (2016). Assessing the severity of wind gusts with lidar. *Remote Sens*, 8, 758. doi:10.3390/rs8090758.
- Frehlich, R.; Meillier, Y.; Jensen, M.L.; Balsley, B.; Sharman, R. (2006). Measurements of boundary layer profiles in an urban environment. *J. Appl. Meteorol. Climatol*, 45, 821–837.
- van der Hooft, E.L.; Schaak, P.; van Engelen, T.G. (2003). *Wind Turbine Control Algorithms*, Energy Research Centre of The Netherlands: Petten, The Netherlands.
- van der Hooft, E.L.; van Engelen, T.G. (2004). Estimated wind speed feed forward control. In: *Proceedings of the European Wind Energy Conference*, 22–25 November 2004, London, UK.
- de Jong, J.J.M.; de Vries, A.C.; Klaasen, W. (1990). Influence of obstacles on the aerodynamic roughness of the Netherlands. *Bound Lay. Meteorol.*, 91, 51–64.
- Sathe, A.R.; Mann, J.; Barlas, T.; Bierbooms, W.A.A.M.; van Bussel, G.J.W. (2013). Influence of atmospheric stability on wind turbine loads. *Wind Energy*, 16, 1013–1032.

3 CHAPTER 1

3.1 First paragraph

3.1.1 Second paragraph

4 CHAPTER 2

5 CHAPTER 3

5.1 First Paragraph

5.1.1 xxx

Appendix A

Table 1. *Table*

Date	Meeting	Location
March 2014	Progress meeting 3	TBD

Table 2.

Table 3.

