SAND2011-8293C

AIAA-2001-0046

PREDICTING DESIGN WIND TURBINE LOADS FROM LIMITED DATA: COMPARING RANDOM PROCESS AND RANDOM PEAK MODELS

LeRoy M. Fitzwater and Steven R. Winterstein

Department of Civil & Environmental Engineering
Stanford University, Stanford, California 94305-4020
leroyf@stanford.edu; steve@ce.stanford.edu

Abstract

This paper considers two distinct topics that arise in reliability-based wind turbine design. First, it illustrates how general probability models can be used to predict long-term design loads from a set of limited-duration, short-term load histories. Second, it considers in detail the precise choice of probability model to be adopted, for both flap and edge bending loads in both parked and operating turbine conditions. In particular, a 3-moment random peak model and a 3- or 4-moment random process model are applied and compared. For a parked turbine, all models are found to be virtually unbiased and to notably reduce uncertainty in estimating extreme loads (e.g., by roughly 50 percent). For an operating turbine, however, only the random peak model is found to retain these beneficial features. This suggests the advantage of the random peak model, which appears to capture the rotating blade behavior sufficiently well to accurately predict extremes.

Introduction

Probabilistic models have gained widespread acceptance and use within a range of engineering disciplines. These models have formed the basis, either explicitly or implicitly, for a number of design codes—especially those of the LRFD (load and resistance factor design) format. Recently developed wind turbine standards (e.g., IEC, 1999) have begun to adopt these code formats, in analogy with long-standing practice in the building and offshore structure communities.

In applying probabilistic models to design wind turbines, a number of practical challenges remain. A first question concerns *how* a particular probability model may be used to satisfy design load checks as specified, for example, in wind turbine standards (e.g., McCoy et al, 1999). In particular, due to the wind turbine's complex dynamic behavior, an analyst may need to rely on a set of limited-duration load histories over a range of wind conditions. These histories may result either from measurements on prototype machines or from numerical simulation. In either case, there is a fundamental question as to how one can proceed from these *short-term* load observations to specification of appropriate *long-term* loads, as required in design codes. The first part of this paper addresses this question, presenting a general methodology to relate the short-term statistics to the desired long-run design load.

A second question that arises concerns the precise details of the probabilistic modelling to be applied; namely, *which* probabilistic model or models are best suited to describe the dynamic behavior of wind turbines. As will be noted below, a number of these have been proposed and applied in the literature. These differ first in which quantity they seek to model; for example, some seek to model the entire load history $x(t)$ as a random process, others seek to model only the local peaks (maxima) of $x(t)$, while still others consider only more global peaks (e.g., 10-minute maxima). Once this choice has been made, various functional forms are available to model the relevant probability distribution at hand. We compare here various random process and random peak models, for cases of both edge and flap bending loads in both operating and parked wind turbine conditions.

Estimating Long-Term Design Loads from Short-Term Histories

In general, LRFD code checks typically compare a nominal load and resistance, L_{nom} and R_{nom}, weighted respectively by factors γ_L and ϕ_R chosen to achieve a desired reliability level:

$$\phi_R R_{nom} \geq \gamma_L L_{nom} \tag{1}$$

The nominal values L_{nom} and R_{nom} are commonly defined somewhat conservatively; e.g., L_{nom}=L_T=the T-year load, defined formally below.

In particular, one proposed wind turbine design check (IEC, 1999) applies the 50-year wind to a parked turbine. This suggests that other checks be made to ensure that this condition is satisfied with L_{nom}=L_{50}, the 50-year *load*, which may not always coincide with the 50-year wind speed. For example, L_{50} may more likely be caused in some cases by turbines operating at lower (but more frequently occurring) wind speeds. We show in this section how one may consistently estimate L_{50}, properly accounting for various wind speeds and their frequencies of occurrence. In the final section of this paper, we show the numerical consequence of underestimating the 50-year load, L_{50}, by considering only the average load under 50-year wind conditions.

As noted above, it is common that the wind turbine analyst may have only limited-duration load histories—formally, observations of M_T, the maximum of the load process, $x(t)$, over a duration T much less than 50 years:

$$M_T = \max_{0 \le t \le T} x(t); \quad T \ll 50 \text{ years} \tag{2}$$

The 50-year load, L_{50}, is then commonly defined as a specific *fractile* of M_T; i.e., a maximum value with a prescribed probability of exceedance:

$$P[M_T > L_{50}] = \frac{T}{50}; \quad T \le 1 \text{ year} \tag{3}$$

Here, $P[\cdot]$ denotes the probability that the bracketed statement occurs. For example, with T=1 year Eq. 3 states that the annual maximum load, $M_{1\ year}$, exceeds L_{50} with probability 1/50=0.02. (Note that we reserve the "L" notation here for deterministic load levels, such as prescribed fractile levels, that arise in design code checking equations. In contrast, M_T denotes the actual maximum load over duration T, which is a random variable.)

Significantly, because of the small probabilities involved, Eq. 3 will return virtually the same L_{50} value for all $T < 1$ year; e.g., by seeking the monthly maximum with exceedance probability .02/12, the daily maximum with exceedance probability .02/365, and so forth. Therefore, in practice one typically reduces T to a duration that can be considered *stationary*; i.e., during which the underlying environmental process (here, the wind speed) can be considered to remain in a statistical steady-state. For wind applications, this may commonly be taken as T=10 minutes, for which Eq. 3 yields L_{50} as

$$P[M_{10\ min} > L_{50}] = \frac{10}{50 \times 365 \times 24 \times 60} = 3.8 \times 10^{-7} \tag{4}$$

By reducing T from 50 years to 10 minutes, we gain the important advantage that the wind speed process remains in a steady state, characterized for example by V, the mean speed during that 10-minute duration. We may then perform a set of steady-state simulations at various mean wind speeds, V, and weight their results by $f(V)$, the long-term probability density of V at the site of interest:

$$P[M_{10\ min} > L] = \int_{all\ V} P[M_{10\ min} > L \mid V] f(V) dV \tag{5}$$

Note that Eq. 5 separates the calculation of L_{50} into the need to provide two separate terms, which respectively describe the turbine (independent of the site) and the wind conditions at the site:

Turbine-specific term: $P[M_{10\ min} > L \mid V]$ denotes the probability that a 10-minute maximum load exceeds a given level L, given a prescribed mean wind speed V. This is commonly known as the *short-term* problem.

Site-specific term: The remaining term on the right side of Eq. 5, $f(V)dV$, defines the fraction of time the wind speed at the site lies between V and $V+dV$. It is common to choose a Rayleigh probability density form for $f(V)$, whose mean depends on site conditions. In general, this wind speed distribution may be found from site-specific data, or specified for design purposes by wind turbine standards (e.g., IEC, 1999; wind site classes I–IV).

In summary, the 50-year load is calculated by first solving the short-term problem—that is, estimating $P[M_{10\ min} > L \mid V]$ across a range of L for various mean speeds V. These results are combined with $f(V)$ through Eq. 5, and L_{50} found as the L value returning the required probability level in Eq. 4.

Note also that Eq. 5 is readily modified if the wind process is instead characterized by another parameter such as turbulence intensity I—replacing V by I—or by a two-dimensional integration if both V and I are deemed to significantly help explain the observed variations in loading. In any case, a common challenge remains to estimate the probability distribution of maximum load; e.g., the 10-minute maximum $M_{10\ min}$. The remainder of this paper considers and applies various models for this purpose.

Probability Models for Extreme Loads and Responses

To estimate the probability distribution of the maximum M_T in time T, one may construct probability models over a number of different time scales. In order of decreasing time scales (and hence increasing use of data), these include the following:

Global Extreme Models: These seek to directly model the "global" (largest) extreme over T, or some slightly lesser duration (e.g., T=1 min or 5 min if not T=10 min directly). The advantage here is that we work most directly with the extreme of interest; i.e., $M_{10\ min}$. The drawback is that we discard all time history values below these global maxima.

Local Extreme/Random Peak Models: These instead model all local maxima of the load history $x(t)$, possibly excluding those that fall beneath some user-defined lower-bound threshold. (This is sometimes referred to as a "peak-over-threshold" model.) Compared with the global extreme models, local extreme models have the advantage of including more data in the fitting process. A potential disadvantage is that some of these data—in particular the lower-amplitude maxima—may come from a different statistical population, which should not be included in extrapolating to large loads. (We show below that this is avoided by appropriate choice of the lower-bound threshold.)

Random Process Models: Finally, these seek to statistically describe the entire time-varying load history, $x(t)$. These contain the largest possible information; i.e., all data points in the digitized history. They may yield little advantage, however, over random peak models if there is little statistically independent information contained in the details of the time history between its peak values.

Note that if we model global extremes directly, we immediately have the desired probability, $P[M_T < L]$, that the maximum value M_T is less than any L. If we instead model all random peaks (here denoted Y_1, Y_2, ...), the corresponding probability $P[M_T < L]$ can be estimated as

$$P[M_T < L] = \{P[Y_i < L]\}^{N_T} \tag{6}$$

in which N_T is the number of peaks (Y_i values) in time duration T. Eq. 6 assumes both that the number of peaks, N_T, is deterministic and that their levels are mutually independent. (Neither assumption is strictly correct, but corrections generally become insignificant as we consider extremes in the upper tails of the load's probability distribution.) Finally, if we instead model the entire process $x(t)$, consistent statistics of M_T require somewhat additional effort. Essentially, one first estimates statistics of local peaks Y_i consistent with the statistics of $x(t)$, and then uses Eq. 6 (see the discussion of functionally transformed Gaussian processes, and their peaks, in Fitzwater and Winterstein, 2000). As noted in that reference, an algorithm (`MaxFits`) has been created to automate this process, permitting the user to select between these three types of models to estimate extreme statistics.

Numerical Results

This section considers how the foregoing probability models can be applied to estimate extreme bending loads on wind turbines. The database we use contains multiple 10-minute simulations of Gaussian wind fields, and corresponding in- and out-of-plane bending moment loads on a specific horizontal axis wind turbine (the Aerodynamics Experiment Phase III turbine; see Madsen et al, 1999 and its associated references). The turbine has a rotor diameter of 10m and a nominal rotor speed of 1.2 Hz. It is a three-bladed turbine with a hub height of 17m.

A total of 100 10-minute simulations have been performed for various choices of the mean wind speed V. These use a general-purpose, commercially available structural analysis code (ADAMS), linked with special-purpose routines to estimate aerodynamic effects (Hansen, 1996). We focus here on three cases:

1. V=14m/s, typical of nominal or "rated" wind conditions;
2. V=20m/s, the maximum or "cut-out" wind speed at which the turbine operates; and
3. V=45m/s, an extreme wind speed (e.g., 50-year level) during which the turbine is parked.

The last case is somewhat analogous to extreme winds on buildings and other stationary structures, and we may expect similar statistical behavior in this wind turbine analysis. The lower-speed cases, however, correspond to operating conditions, in which the turbine blades rotationally sample the stationary wind field. Also notable here are the systematic effects of gravity on in-plane bending: a strong sinusoidal trend is induced at the turbine operating speed. We investigate here whether various probabilistic response models can remain accurate in the face of these special features that wind turbines exhibit.

In particular, we study here the behavior of two different types of probabilistic models: (1) *Hermite* models of the load as a random process, and (2) *quadratic Weibull* models of random load peaks (over a specified threshold). The Hermite model generally utilizes four

statistical moments of $x(t)$ (Winterstein, 1988), although a simplified three-moment version can be used in some special cases of limited nonlinearity (Jha and Winterstein, 1997; Jha and Winterstein, 2000; Madsen et al, 1999). The quadratic Weibull model is based on the first three statistical moments of the peak values, Y_i (Lange and Winterstein, 1996; Ronold et al, 1999; Winterstein and Kashef, 1999). Note that the `MaxFits` routine implements both the quadratic Weibull and Hermite models, as well as a number of others.

Results 1: Sample Time Histories

Figure 1 shows simulated wind and load time histories from one 10-minute simulation. In particular, the histories shown are brief, 5-second portions of the wind and load histories during which the wind input is maximized. (This maximum wind episode does not generally produce the maximum bending loads.)

To identify peaks from the response histories, we define a peak here as the largest value of the history between successive upcrossings of its mean level*. Figure 1 shows the mean levels of each history by horizontal lines, and the circled response points indicate the set of peaks that are obtained. The out-of-plane (flap) bending loads are found here to roughly follow the wind speed process, although additional high-frequency content is observed. Note also that our definition of peaks (largest response per upcrossing of the mean) serves to filter out many of these high-frequency response oscillations. The edge bending loads are of less interest in this case, showing small oscillations about the mean load.

Figure 2 shows similar simulated wind and load time histories, now from a wind speed V=20m/s during which the turbine is operating. Now the effect of gravity is clearly seen in the edge bending history, which shows a strong sinusoidal component at the operating speed of roughly 1.2 Hz. The flap bending history also shows systematic variations at this frequency, although it is combined with significantly larger high-frequency content here than in the edgewise case. Again, our peak identification method removes some of this high-frequency effect. Note in the edgewise case, however, that a somewhat anomalous effect can arise. While only one "large amplitude" peak is usually found per blade revolution, other "secondary", near-zero peaks are sometimes also identified. This arises from the high-frequency small-amplitude oscillations shown by the edgewise loads about their mean level. The resulting distribution of all peaks is found in such cases to be multi-modal; i.e., to possess a probability density function with several distinct regions of relatively high probability ("modes"). Because our models are unimodal—i.e., designed to be fit to the single most important probability "mode"—we shall find it useful in these edgewise cases to pass a higher threshold (above the mean) to exclude these secondary peaks. We shall return to this issue below.

Finally, recall that to estimate the distribution of the *largest* peak, it is common to assume that successive peaks are mutually independent. This is the assumption inherent in our current implementation of `MaxFits` (see, for example, Eq. 6). To test this assumption, the correlation coefficient ρ between adjacent peaks Y_k and Y_{k+1} has been calculated for various cases. Typical ρ values are effectively negligible; for example, ρ=0.21 (flap bending, V=45m/s), ρ=0.28 (edge bending, V=45m/s), and ρ=0.15 (flap bending, V=20m/s). These values, and corresponding scatterplots of Y_k vs Y_{k+1} (e.g., Figure 3), confirm that the assumption of independence should not induce large modelling errors in this application. This conclusion may differ in other applications; for example, the lightly damped slow-drift response of some moored marine structures.

Results 2: Observed vs Predicted Distributions of Peaks

We now test the ability of a three-moment, quadratic Weibull distribution to accurately model the simulated response peaks across various wind conditions. For illustration purposes, we again show results for the first (of the 100) 10-minute simulations.

We again consider first the parked turbine (V=45m/s), whose statistical behavior may be expected to be most well-behaved. Figure 4 shows the cumulative probability distribution function $F_Y(y)$=$P[Y \le y]$ of all peaks, as estimated directly from the data. Specifically, for both flap and edge cases, the peaks y_i are first ordered so that $y_1 \le y_2 \le ... \le y_n$, and associated with the cumulative probabilities p_i=$F_Y(y_i)$=$i/(n+1)$. Results are plotted on a distorted "Weibull" scale, which plots y not versus $F_Y(y)$ but rather versus $-\ln[1-F_Y(y)]$. The results, when viewed on log-log scale, should appear as a straight line if the data follow a Weibull probability distribution model.

The data here show slightly positive curvature on this Weibull scale. This suggests the value of the quadratic Weibull model, which yields a quadratically varying distribution when plotted on the Weibull scale of Figure 4. This quadratic model is shown here to ac-

*Many alternative strategies can be used to identify peaks; e.g., the largest value per blade revolution. Note, however, that for these edge load cases we find an optimal threshold to exclude somewhat more load cycles; e.g., we retain roughly 1 peak per 2 blade revolutions.

curately follow both the flap and edge load data in this case.

Figure 5 shows similar Weibull scale plots of flap and edge loads in the V=20m/s case, during which the turbine is rotating. While the distribution of flap load peaks remains smooth, the distribution of edge load peaks shows a sharp change in behavior, with a "corner" located at roughly y=1. This is a consequence of the bimodal character of the edge load peaks, as discussed earlier. No smooth, single-moded distribution model can capture both the large, one-per-revolution primary peaks and the small-amplitude, secondary peaks. For both ultimate and fatigue load modelling purposes, however, these secondary peaks are of little consequence. We therefore seek to model the *shifted* peaks, $Y-1.5$; i.e., we remove all peaks below 1.5, and report the shifted values y_i'=$y_i-1.5$ of the remaining peaks. The shifting is used to conform with quadratic Weibull models, which generally assigns probability to all outcomes $y'\geq 0$. Figure 6 shows the quadratic Weibull model to accurately follow the shifted edge loads, $Y-1.5$. (Note that the optimal choice of shift parameter may require some trial and error; e.g., comparing goodness-of-fit measures. This is a topic of ongoing study. Note also that in using these models to predict extremes, the shift value must eventually be reinstated, to report loads in units consistent with their input values. This is done automatically in the `MaxFits` routine).

Results 3: Estimating 10-Minute Mean Maxima

Finally, we show predicted statistics of $M_{10\ min}$, the maximum 10-minute load. In particular, we seek here to estimate m, the mean value of $M_{10\ min}$ to be expected in an arbitrary 10-minute period. As shown below, this mean value is critical in fitting the probability distribution of $M_{10\ min}$, for use in the long-term load calculation of Eq. 5.

A simple, "raw" estimate of m can be found by averaging the 100 observed maxima, M_i, from each of the 10-minute simulations:

$$\overline{M}=\frac{1}{100}\sum_{i=1}^{100}M_i \tag{7}$$

Alternatively, we can estimate m by fitting one of the foregoing models; e.g, a quadratic Weibull model to all response peaks (perhaps above a shifted level). Here we fit such models separately to each of the 100 simulations. Denoting μ_i as the estimated value of m from simulation i (i=1,...,100), we may form an analogous average of these estimates:

$$\overline{\mu}=\frac{1}{100}\sum_{i=1}^{100}\mu_i \tag{8}$$

One advantage of the simple, "raw" estimate $\overline{M}$ is that it is always "unbiased"; i.e., correct on average. A potential disadvantage is that because it is based on only the single observed maximum in each 10-minute history, it may show considerable variability. By instead fitting probability models to form estimates μ, we hope to achieve results that (1) remain nearly unbiased and (2) show reduced scatter (specifically, standard deviation) compared with the raw estimate $\overline{M}$. To quantify these effects we define two factors: a bias factor, defined as

$$\text{Bias (B)}\ =\frac{\overline{\mu}}{\overline{M}} \tag{9}$$

and a sigma reduction factor, defined as

$$\text{Sigma Reduction (SR)}\ =\frac{\sigma_\mu}{\sigma_M} \tag{10}$$

We hope to achieve bias factors of nearly unity, and sigma reductions far less than unity. Again, our hope for sigma reduction lies with the fact that each estimated μ_i uses more of the simulation history—specifically, each peak-over-threshold value—than the raw estimate $\overline{M}$, which uses only the single maximum M_i from each 10-minute simulation.

Figures 7 and 8 show bias and SR factors, respectively, for the parked turbine (V=45m/s). Three probability models are fit: a 3-moment quadratic Weibull model ("Peak Unshifted"), and both 3- and 4-moment Hermite models of the complete random response process $x(t)$. (The four moment model has become the option of choice for general use. The three-moment simplification has been used in some mildly nonlinear wave applications, and has been recently derived independently for wind turbine applications (Madsen et al, 1999)). Note that all models yield roughly unbiased results (B near 1.0). The 3-moment models generally achieve a sigma reduction of 0.5 or less. As might be expected, inclusion of the 4th moment, with its attendant uncertainty, leads to higher values of σ_μ and hence SR.

Figures 9–10 show analogous bias and SR factors for V=20m/s, one of the operating wind speed conditions. Here the random process (Hermite) models, which

are intended to model rather general stochastic behavior, fail to accurately capture the rotating nature of the blade response. Biases of about 10% are found from conventional (4-moment) Hermite models, with considerably larger biases produced by the simpler 3-moment Hermite models.

In contrast, the quadratic Weibull ("Peak") models remain essentially unbiased in all cases. For cases of edge loads, models have been fit both to the original data y_i ("Unshifted") and the shifted data $y_i - 1.5$ ("Shifted"). For this particular choice of duration (T=10-minute maxima), even the unshifted models appear reasonably accurate. Over longer durations, however, estimates become increasingly tail-sensitive, and the use of the shift has been found more beneficial in avoiding bias. This is reflected in Figure 11, which shows the benefit of including a shift when predicting maxima over a range of T=10–1000 minutes. The shifted predictions also retain the roughly 50% sigma reduction, as shown by the confidence bands in Figure 12. In all operating and parked conditions, sigma reductions for these peak models have been found to remain at roughly 0.5 or less.

Note also that when averaging results over N simulations, the standard deviation of an estimated parameter decreases like $\sigma/\sqrt{N}$. Hence, the 50% sigma reduction shown by the quadratic Weibull fit permits a four-fold decrease in the number of simulations. For example, fitting a quadratic Weibull model to N=1 10-minute simulation is roughly equivalent to using the "raw" 10-minute maxima from N=4 simulations.

Example of Long-Term Load Analysis

Finally, we show how the foregoing models of 10-minute maxima, $M_{10\ min}$, can be used to predict long-term loads as in Eq. 5. The 10-minute mean wind speed, V, is assumed to have Rayleigh probability distribution with mean V_{ave}:

$$P[V \geq v] = G_V(v) = \exp\left[-\frac{\pi}{4}\left(\frac{v}{V_{ave}}\right)^2\right] \qquad (11)$$

Setting $G_V(v)$ to the 50-year exceedance probability (Eq. 4) yields V_{50}=4.34V_{ave}. Here we assume V_{50}=45 m/s; i.e., the parked turbine simulations correspond to 50-year wind conditions. This implies V_{ave}=45/4.34=10.4 m/s, roughly corresponding to a class I wind site (IEC, 1999).

For load modelling purposes, we seek to estimate $m(V)$ and $\sigma(V)$, the mean and standard deviation of $M_{10\ min}$, as functions of V. To illustrate, we assume here simple linear relations:

$$m(V) = c \cdot \frac{V}{V_{50}}; \quad \sigma(V) = d \cdot \frac{V}{V_{50}} \qquad (12)$$

	c [kN-m]	d [kN-m]
$V \leq 20$m/s	11.3	1.13
$V > 20$m/s	20.0	1.63

Table 1. Values of coefficients c and d used in Eq. 12 to fit flap load moments as functions of the mean wind speed, V.

These results have been fit separately to the operating and parked turbine simulations ($V \leq V_{cut\ out}$=20 m/s and $V >$20 m/s). Table 1 shows corresponding c and d coefficients for the flap load case. Figure 13 compares the observed and fitted moments in this case. (Of course, in practice these fits should use results over a greater set of wind conditions, and different functional forms may emerge. Our results here are intended only for illustration purposes.)

An important simplification arises if the mean load $m(V)$ grows steadily with V, and the load variability $\sigma(V)$ is small compared with the variability in V. In this case we may estimate the 50-year load, L_{50}, as $m(V_{50})$, the mean load in 50-year wind conditions. This is essentially the current design load check for parked turbines (load case 6.1 in IEC, 1999). Formally, this will be exact in the "deterministic" load limit, i.e., as $\sigma(V) \to 0$. More generally it will be somewhat unconservative. In offshore structure design, for example, this unconservatism is noted and commonly adjusted by (1) inflating the environmental variable (here, the wind speed or associated turbulence); or (2) inflating the fractile (from p=.5 to p=.80–.90) at which the load is evaluated. These inflation procedures are basically empirical, and have been calibrated with respect to long-term probability analysis (as in Eq. 5) across many cases (e.g., Haver et al, 1998; Winterstein and Engebretsen, 1998).

Figure 14 shows the resulting long-term distribution of 10-minute flap loads. The "random" load case uses Eq. 5, with the discretized probability $f(V)\Delta V$ taken as the difference between cumulative probability values (Eq. 11) evaluated at the lower and upper v_i values for each bin. The remaining term in Eq. 5, $P[M_{10\ min} > L|V]$, assumes that $M_{10\ min}$ has a Gumbel distribution with mean and standard deviation as given in Eq. 12. The corresponding "deterministic" result sets $\sigma(V)$=0, so that

$P[M_{10\ min} > L]$ is $P[m(V) > L]$, which can be evaluated analytically with Eq. 11. (It has also been evaluated here by the numerical integration with $\sigma(V)$=0, to verify the accuracy of the discretization.)

For example, in the deterministic case the 50-year flap load is predicted as $m(V_{50})$=20 [kN-m], as seen by evaluating $m(V)$ in Figure 13 at V_{50}=45m/s. In contrast, including realistic load randomness leads to L_{50}=22.7, an increase of 13.5%. This is not inconsistent with offshore structural responses, which show roughly 10–20% increases due to load randomness.

Finally, note that the long-term results in Figure 14 include randomness in the loading, but ignore uncertainty in load moments due to limited data. Figure 15 shows upper-bound 95% confidence results for this long-term load distribution. These estimate the moments $m(V)$ and $\sigma(V)$ by the raw sample mean and variance, using either N=4 or N=8 10-minute simulations. (Recall that this is roughly equivalent to using a quadratic Weibull fit, applied to N=1 or N=2 simulations.) Based on the variability of sample means and variances, we model $m(V)$ and $\sigma(V)$ as

$$m(V) = \hat{m}(V) + \frac{\hat{\sigma}(V)}{\sqrt{N}} \cdot U; \quad \sigma(V)^2 = \hat{\sigma}(V)^2 \cdot \frac{N-1}{\chi^2_{N-1}} \tag{13}$$

Here $\hat{m}(V)$ and $\hat{\sigma}(V)$ are our best estimates, as found from Eq. 12, while U and χ^2_{N-1} are standard normal and chi-square variables, the latter with N-1 degrees of freedom. Results here are based on 100 simulations, whose resulting probabilities $P[M_{10\ min} > L]$ have been sorted (for various L) and the 95% fractile reported. Note the marked consequence of limited data: with N=4 the 50-year load estimate should be increased by roughly 50% to achieve 95% confidence. This effect can be reduced of course through additional data: with N=8 simulations the 50-year load increases by about 20%. Recall also that we may interpret these N=4 and N=8 cases as being roughly equivalent to using the quadratic Weibull model with only N=1 or 2 simulations. This again shows the benefit of applying such a fit, to use data more efficiently and hence reduce (for fixed N) our upper-fractile load distribution estimates.

Summary

This paper has shown how, through the use of probability models, short-term loads data can be used to predict long-term design values. Regarding the specific choice among these models, it has demonstrated the use of both random process and random peak models of wind turbine loads. In particular, it has applied 3-moment random peak models (quadratic Weibull), and 3- and 4-moment random process models (Hermite). Both the quadratic Weibull and (4-moment) Hermite models are available within `MaxFits`. Sample input and output file have been created (Fitzwater and Winterstein, 2000) to illustrate how this routine can be used to derive some of the results shown here.

For a parked wind turbine experiencing 50-year winds, all models have been shown to be nearly unbiased (Figure 7) and to achieve a significant reduction in our uncertainty (Figure 8) in estimating the mean 10-minute maximum. For rotating blades during operation (at lower wind speeds), the random process models can show notable bias: roughly 10% for the 4-moment models, and appreciably more if only 3 moments are used (Figures 9–10). In contrast, the random peak models remain consistently accurate, and consistently beneficial (i.e., in reducing uncertainty) in all cases. This suggests that by modelling not the entire time history but rather its set of peaks, enough information about the rotating nature of the load process is retained to permit accurate estimates of extreme behavior.

Acknowledgements

We wish to thank Kirk Pierce, Marshall Buhl, and their associates for making available the simulation results of the Aerodynamics Experiment Phase III turbine. More generally, we gratefully acknowledge a series of useful—and ongoing—discussions with various members of the Statistical Loads Extrapolation (SLEx) program. This is an informal working group, with representatives from NREL, Sandia, consultants, industry, and universities. Through interactions with this group, we hope to make reliability methods more accessible to the wind turbine community, and more directly tailored to their specific needs. Finally, this work is supported by the Wind Energy Program of the U.S. Department of Energy, through Sandia National Laboratories.

References

Fitzwater, L.M. and S.R. Winterstein (2000). *Estimation of extremes from limited time histories: The routine MaxFits with wind turbine examples*, Rept. RMS–39, Reliability of Marine Structures Program, Dept. of Civ. & Environ. Eng., Stanford University. To be issued as a technical report by Sandia National Laboratories.

Hansen, A.C. (1996). *Users guide to the wind turbine*

dynamics computer programs YawDyn and AeroDyn for ADAMS, Mech. Eng. Dept., Univ. of Utah.

Haver, S., G. Sagli and T.M. Gran (1998). Long term response analysis of fixed and floating structures. *Proc., Wave'98—Ocean Wave Kinematics, Dynamics and Loads on Structures*, International OTRC Symposium, Apr. 30–May 1, 1998.

IEC, 1400-1 (1999). *Wind turbine generator systems, part 1: safety requirements*, International electrotechnical commission, Report IEC 61400-1, ed. 2.

Jha, A.K. and S.R. Winterstein (1997). *Nonlinear random ocean waves: prediction and comparison with data*, Rept. RMS–24, Reliability of Marine Structures Program, Dept. of Civ. Eng., Stanford University.

Jha, A.K. and S.R. Winterstein (2000). Nonlinear random ocean waves: prediction and comparison with data. *Proc., 19th Intl. Offshore Mech. Arctic Eng. Symp.,* ASME, Paper No. OMAE 00–6125.

Lange, C.H. and S.R. Winterstein (1996). Fatigue design of wind turbine blades: load and resistance factors from limited data. *Proc., Wind Energy 1996,* ASME, SED **17**, 93–101.

Madsen, P.H., K. Pierce and M. Buhl (1999). Predicting ultimate loads for wind turbine design. *Proc., 1999 ASME Wind Energy Symposium*, 37th AIAA Aero. Sci. Mtg., 355–364.

McCoy, T.J., D.J. Malcolm and D.A. Griffin (1999). An approach to the development of turbine loads in accordance with IEC 1400–1 and ISO 2394. *Proc., 1999 ASME Wind Energy Symposium*, 37th AIAA Aero. Sci. Mtg., 1–9.

Ronold, K.O., J. Wedel-Heinen, and C.J. Christensen (1999). Reliability-based fatigue design of wind-turbine rotor blades. *Engineering Structures*, **21**, 1101–1114.

Winterstein, S.R. (1988). Nonlinear vibration models for extremes and fatigue. *J. Engrg. Mech.,* ASCE, **114**(10), 1772–1790.

Winterstein, S.R. and K. Engebretsen (1998). Reliability-based prediction of design loads and responses for floating ocean structures. *Proc., 17th Intl. Offshore Mech. Arctic Eng. Symp.,* ASME.

Winterstein, S.R. and T. Kashef (1999). Moment-based load and response models with wind engineering applications. *Proc., 1999 ASME Wind Energy Symposium*, 37th AIAA Aero. Sci. Mtg., 346–354.

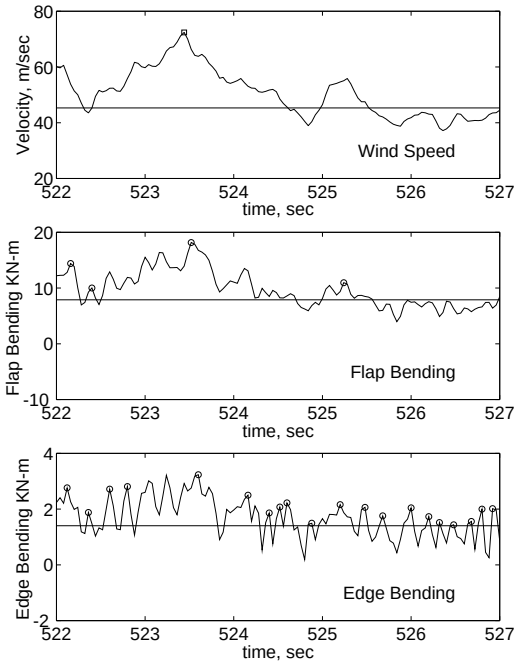

Figure 1. Simulated wind and blade loads; V=45m/sec.

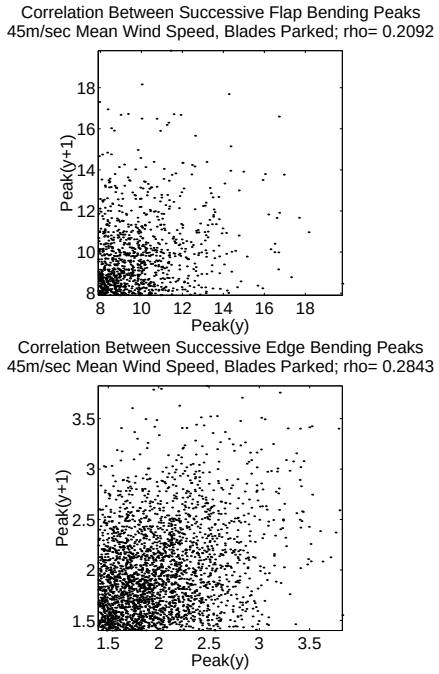

Figure 3. Correlation between successive peaks; V=45m/sec.

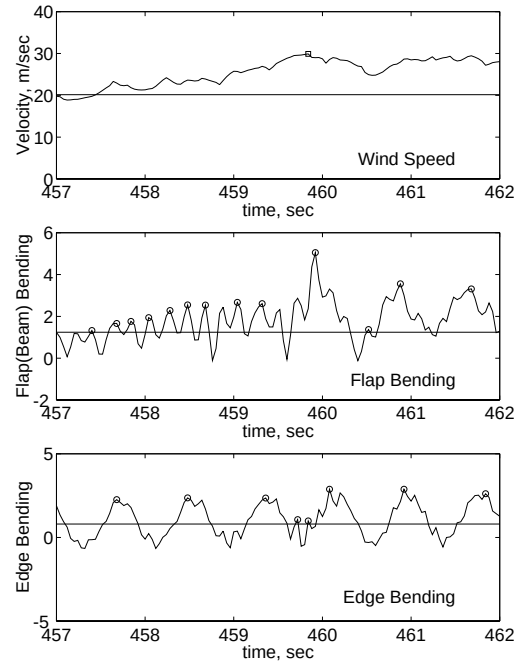

Figure 2. Simulated wind and blade loads; V=20m/sec.

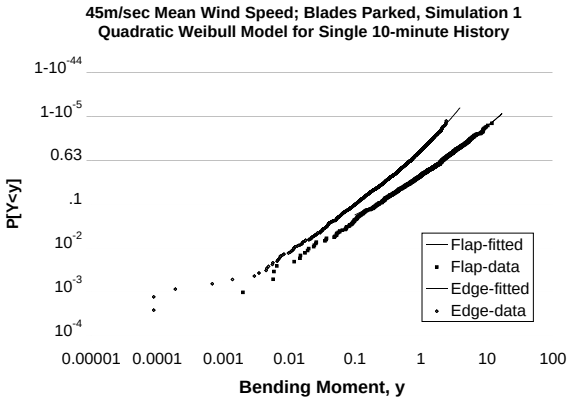

Figure 4. Probability distribution of response peaks; V=45m/sec.

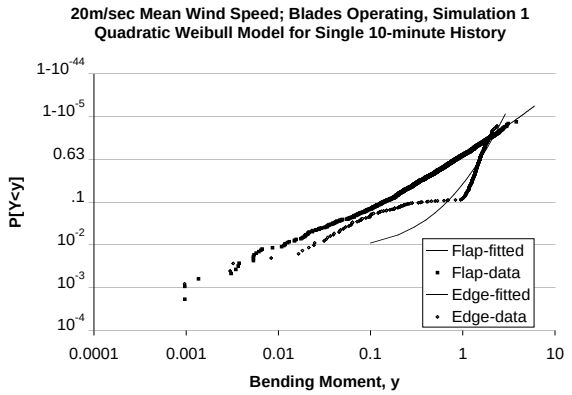

Figure 5. Probability distribution of response peaks; V=20m/sec.

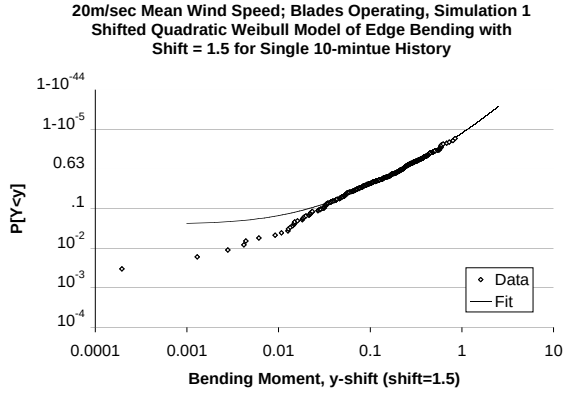

Figure 6. Probability distribution of *shifted* response peaks above 1.5; V=20m/sec.

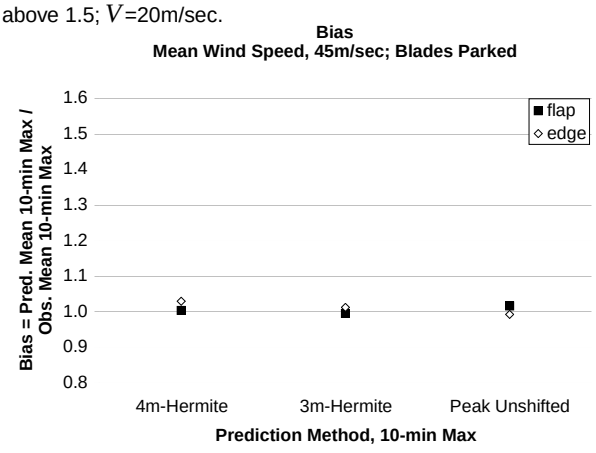

Figure 7. Bias = $\overline{\mu}/\overline{M}$, where $\overline{\mu}$ is the average estimate of the mean 10-min maximum over the 100 simulations. $\overline{M}$ is the average of the observed 10-min maxima. The wind speed is V=45m/sec.

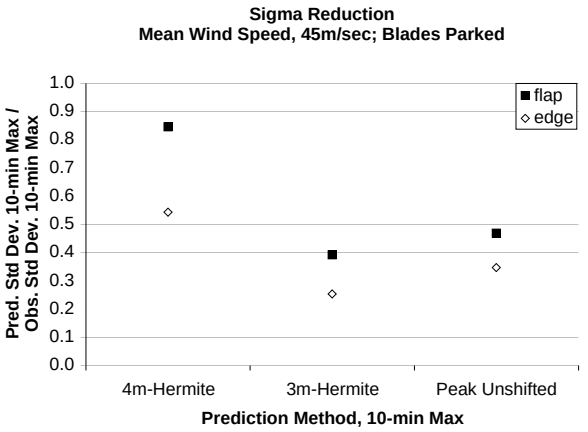

Figure 8. Sigma reduction σ_μ/σ_M between estimated and observed 10-min maxes; V=45m/sec.

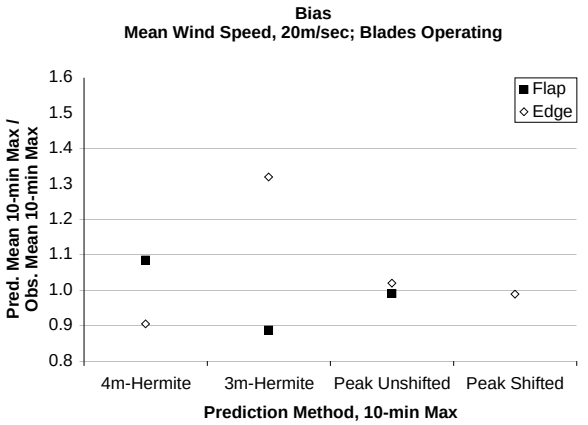

Figure 9. Bias = $\overline{\mu}/\overline{M}$, where $\overline{\mu}$ is the average estimate of the mean 10-min maximum over the 100 simulations. $\overline{M}$ is the average of the observed 10-min maxima. The wind speed is V=20m/sec.

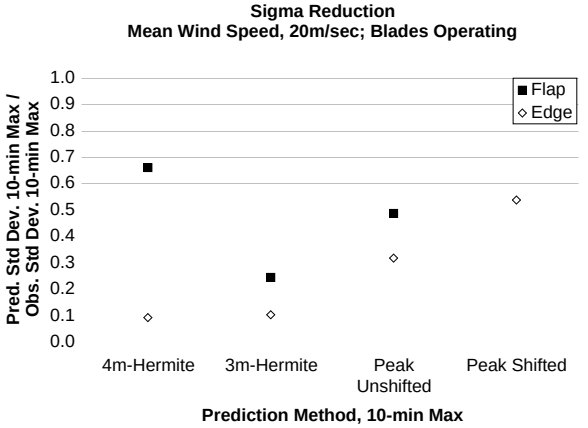

Figure 10. Sigma reduction σ_μ/σ_M between estimated and observed 10-min maxes; $\overline{V}$=20m/sec.

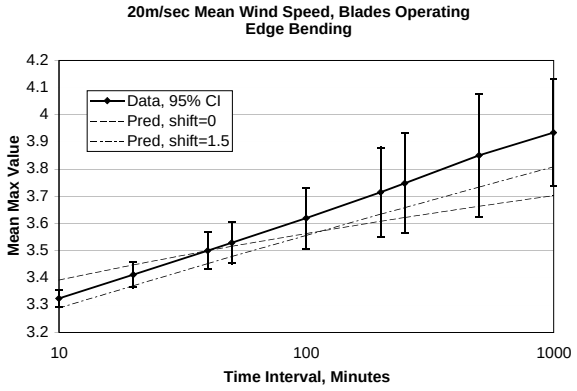

Figure 11. Estimated mean maxima over various time intervals; edge bending V=20m/sec.

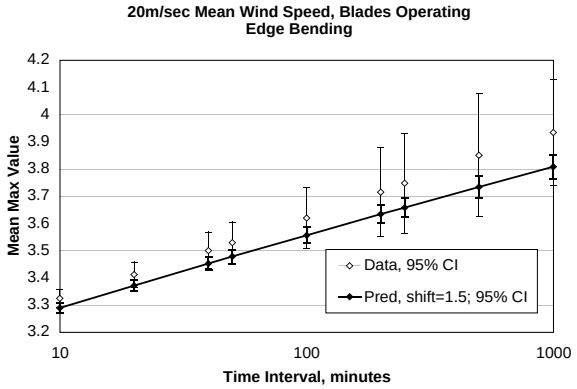

Figure 12. Estimated mean maxima over various time intervals; edge bending V=20m/sec.

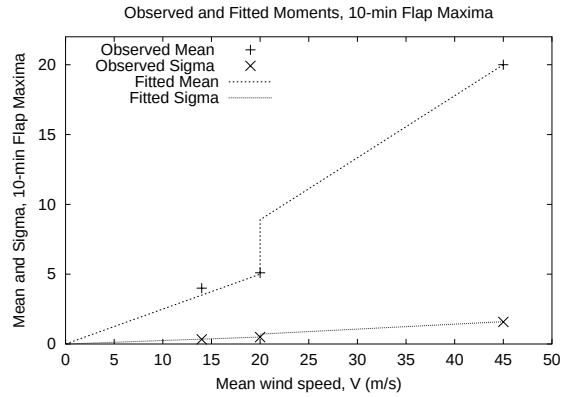

Figure 13. Mean and standard deviation of 10-minute maxima; observed and fitted results.

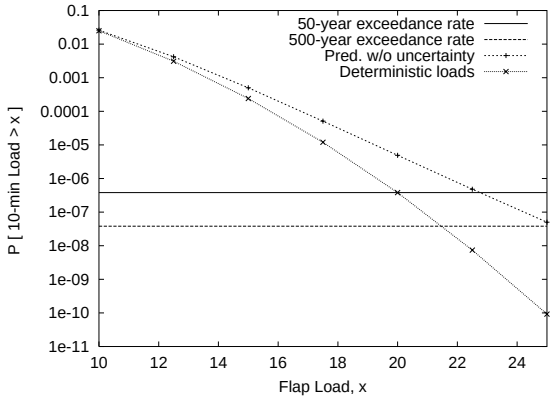

Figure 14. Long-term distribution, including and excluding load randomness.

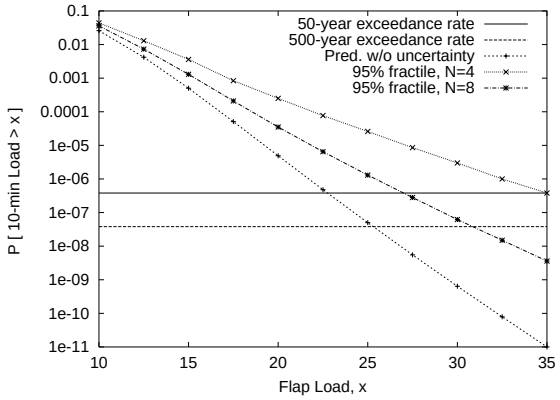

Figure 15. Long-term distribution, uncertainty using $N = 4$ or 8.