**FATIGUE RELIABILITY OF WIND TURBINE FLEETS:
THE EFFECT OF UNCERTAINTY ON PROJECTED COSTS***

Paul S. Veers
Sandia National Laboratories
Wind Energy Technology Department
MS - 0708
Albuquerque, NM 87185-0708
USA
Phone (505)844-5522
Fax (505)845-9500
psveers@sandia.gov

ABSTRACT

The cost of repairing or replacing failed components depends on the number and timing of failures. Although the total probability of individual component failure is sometimes interpreted as the percentage of components likely to fail, this perception is often far from correct. Different amounts of common versus independent uncertainty can cause different numbers of components to be at risk of failure. The FAROW tool for fatigue and reliability analysis of wind turbines makes it possible for the first time to conduct a detailed economic analysis of the effects of uncertainty on fleet costs. By dividing the uncertainty into common and independent parts, the percentage of components expected to fail in each year of operation is estimated. Costs are assigned to the failures and the yearly costs and present values are computed. If replacement cost is simply a constant multiple of the number of failures, the average, or expected cost is the same as would be calculated by multiplying by the probability of individual component failure. However, more complicated cost models require a break down of how many components are likely to fail. This break

* This work was supported by the United States Department of Energy under Contract DE-AC04-94AL85000.

down enables the calculation of costs associated with various probability of occurrence levels, illustrating the variability in projected costs. Estimating how the numbers of components expected to fail evolves over time is also useful in calculating the present value of projected costs and in understanding the nature of the financial risk.

INTRODUCTION

The fatigue life of many wind turbine components is susceptible to large uncertainties for two reasons. First, the fatigue resistance of all materials has a large amount of inherently random scatter. That is, lifetimes of identical materials stressed identically often differ by factors of ten or even hundreds. Second, small changes in loading will lead to large changes in lifetime. This sensitivity exacerbates the problem of not knowing the loadings perfectly. Small uncertainties in the loadings lead to large uncertainties in component lifetimes. The sum of these two effects is to create a wide range of possible lifetimes for fatigue-susceptible wind-turbine components.

An explanation of how fatigue life can be calculated for wind turbines can be found in several references (Sutherland, et al., 1994; Spera, 1994; Jackson, 1992; and Veers, 1989 and 1982). Component fatigue life is usually calculated using the best estimates of uncertain load and resistance quantities and applying reasonable safety factors. An equivalent approach is to use a selected confidence level for each input quantity.

A better measure of design adequacy may be obtained by calculating the probability of component failure before a specified target lifetime (Veers, et al., 1994). The probabilistic approach uses the distribution of possible values for uncertain inputs and estimated the probability of achieving a desired lifetime for a component. It is usually more meaningful to know the probability of achieving an desired lifetime than simply knowing how much greater the lifetime will be on average.

The economic impact of uncertainty in fatigue life is addressed in this paper by calculating the expected cost of replacing failed components in a fleet of turbines. The probability of different percentages of components failing is first estimated for a range of target lifetimes. Costs are calculated based on the probability that different numbers of components will fail in each year of operation and the expense of replacement. The cost in each operating year is then known, and the present value can also be estimated. Both average costs and costs associated with different percentile levels of probability of occurrence are calculated. Thus, the nature of the risk to the fleet is both illustrated and quantified for use in making financial decisions.

PROBABILITY OF FAILURE

The probability of wind turbine components failing to meet a target lifetime can be evaluated with a software tool called FAROW, for Fatigue And Reliability Of Wind turbines (Veers, et al., 1994). FAROW uses structural reliability methods to evaluate the probability of premature failure in the presence of multiple uncertain inputs. It is specifically tailored to the wind turbine fatigue problem and does all the probability calculations internally, leaving the user to focus attention on the still formidable task of defining the input distribution functions. Table 1 lists the inputs to a fatigue lifetime calculation that are typically uncertain. Each input can be defined as either a constant or a random variable with a probability distribution taken from a library of functional forms.

FAROW calculates the median lifetime of the component, the probability of failing before the user-specified target lifetime, and importance factors, which indicate how much each uncertain input contributes to the probability of failure. The output of most interest here is the probability of a component failing in less than a specific target lifetime. In the past, it was impossible to assess the cost penalty associated with uncertainty. FAROW now makes it possible to conduct a detailed economic analysis of the effects of uncertainty in component repair or replacement due to fatigue.

Common and Independent Uncertainties

When the probability of a component failing in less than Y years is calculated by FAROW to be X%, one often hears the interpretation that “you can expect X out of 100 components to fail in the first Y years of operation.” Unfortunately, this very simple and useful way to think is usually wrong. It would be correct if all the uncertainties in the inputs are completely *independent* from component to component. However, much of the uncertainty does not lie in the randomness of an input quantity from component to component. Rather, the quantity has some value that varies quite little from component to component, but the exact value of the quantity is simply not known. This uncertainty is *common* (perfectly correlated) between all the machines in the fleet. If *all* of the uncertainty is common between all the components, the correct interpretation of the above statement would be that either none or all of the components will fail, and the probability of all of them failing is X%. Real life is never so simple as to fit into either limiting category, but contains uncertainty of both the common and independent varieties.

Completely separating the uncertainty in component fatigue life into common and independent sources is a virtually impossible task, or at best very difficult. However, material fatigue properties have such a large and inherently independent variability that they can be used to approximate all the independent uncertainty. Figure 1 shows fatigue test results for identical specimens, plotted as effective, alternating-stress amplitude versus number of cycles to failure, i.e., a stress-life, or S-N, plot. Notice that identical material specimens tested at the same stress level can have lifetimes that differ by a factor of ten or more. This is the norm and not the exception with fatigue properties. A typical value for the standard deviation of the cycles-to-failure is 60% of the mean value (ASCE, 1982). There may also be common material property uncertainty due to manufacturing processes and material lot differences. These, however, are assumed to be small relative to the inherent randomness of the material property.

Material Fatigue Properties

Material S-N, or stress-life, properties are often defined using the functional form of

$$N = CS^{-b} \quad (1)$$

where N is the number of constant amplitude (S) cycles to failure and C and b are material properties. The coefficient C gives the location of the S-N curve and the exponent b gives the slope. This single curve description is usually sufficient over the range of interest for a fatigue life estimate.

The randomness in material S-N properties can be described in FAROW by defining C and b as correlated random variables that reflect the statistics of the fatigue-test data. However, the nature of the data is most simply modeled by assuming the slope is a constant and using a single random variable, the coefficient C of the S-N curve, to describe the uncertainty in lifetime given loading.

The S-N coefficient can also be defined as a constant representing a given “confidence level” or survival rate. Figure 1 shows curves for four survival rates: 50% (Least Squares Curve Fit), 90%, 95%, and 99%. When the material properties are thus entered in FAROW as constants describing an S-N curve at a selected survival rate, and all the other inputs listed in Table 1 are input as usual, the calculated probability of failure reflects the probability that *more than* the selected percentage of components will have failed *before* the target lifetime.

All of the inputs other than the S-N properties, although possessing some independent randomness from machine to machine, are most likely dominated by common sources (i.e., values that are common to all the components, but not known with certainty). Therefore, the material property is chosen to represent all the independent uncertainty and the rest of the inputs are assumed to be entirely common between components. This assumption should be accurate in most applications, and greatly simplifies the analysis presented here, but is not necessary for the application of the procedure.

The result is that we can use FAROW to estimate the probability of achieving fleet-wide survival rates, specified by the S-N curve survival rate, at any designated target lifetime. The process is summarized here:

1. Set the material S-N coefficient to reflect a specified percentage-survival level and input it into FAROW as a constant.
2. Input all the other uncertain inputs (the *common* ones) to FAROW as random variables.
3. Run the FAROW analysis for the probability of failure at several target lifetimes. *Failure* here means that more than the specified percentage of failures will have occurred before the target lifetime.
4. Rerun the analysis for all percentage-survival levels of interest.

COST ESTIMATES

Having calculated the probability of different percentages of the components in a fleet of turbines failing at different target lifetimes, it is a simple matter to apply a cost to the expected fleet-wide repair or replacement. First, calculate the probability that the number of components failing will be in a specific range. For example, the probability of between 1% and 2% failing is the probability of more than 1% failing minus the probability of more than 2% failing. Then multiply the replacement cost by the numbers of components failing (1.5% in the previous example) and weight it by the probability of that occurrence. This yields the *expected* cost in each range of percentages of fleet failures. By breaking the costs down this way, the source of financial risk is exposed, i.e., whether the risk is from a high probability of small numbers of components failing or from a low probability of many components failing. The total expected, or average, cost of fleet maintenance is estimated by summing over all the possible percentages of components failing. Costs can also be broken down into the year (or range of years) in which they are accrued. For example, the probability of failing in an 8-10 year range is the probability of failing in less than 10 years minus the probability of failing in less than 8 years.

Of course costs accrued in the future have a different value in the present, i.e., “present value,” than an equal current expense. In an economic analysis, the present value is obtained by “discounting” the future expense. The present value, P , of future costs can be calculated using

$$P = \sum_i E_i (1-d)^i \quad (2)$$

where E_i is the expected expense in the i^{th} year and d is the “discount rate.”

The discount rate is used to reflect not only inflation, but the cost of capital and various other premiums. It is not the intent here to specify the correct discount rate. Instead the estimates of present value will be illustrated over a range of discount rates. The details of this cost analysis process may be best illustrated in the following example.

EXAMPLE

The process of calculating the economic effect, or financial risk, of uncertainty may be illustrated with an example taken directly from the FAROW User’s Manual (Veers, et al., 1994). The example represents a situation in which extensive prototype testing has reduced the uncertainty in the machine response to the environment about as far as possible. This relatively low uncertainty case has a median lifetime of 300 years, while the probability of the component failing in less than 20 years is 3%.

Descriptions of the input quantities reflecting the appropriate degree of randomness and uncertainty for these examples can be found in the FAROW User’s Manual. The exact inputs are not important to the topic of this article. Rather, they can be summarized using the importance factors, which reflect the contribution to the probability of failure due to each of the random inputs. Figure 2 shows the importance factors for this case lumped into three areas: wind speed, stress response and material property inputs (see also Table 1). Material properties include the inherent randomness in fatigue properties, and represent all of

the independent uncertainty in this example. The wind-speed category describes the long-term wind-speed distribution. The stress-response category includes such quantities as stress-concentration factors, mean-stress levels, stress dependence on wind speed, and load-spectrum shape. As stated above, all the latter two categories are designated as common sources of uncertainty.

The probability that any individual component will fail is calculated as usual, but we would also like to know how many components are likely to fail and at what time. The four steps listed earlier are followed to accomplish this task. Establishing a rigorous zero failure rate is somewhat problematic so a failure rate of less than 0.1% is interpreted here as no failures.

Probability of Component Failure

If all median properties are used to calculate the fatigue life in this example, the life of this component (a blade joint) is estimated at 300 years. Of course, no designer worth his or her salt would ever design with median properties. Some substantial factors of safety would be applied. Here, we calculate the probability that a component will last for a predetermined period of time using the FAROW software. As stated above, this probability of failure in a 20-year lifetime for an individual component has been calculated at 3%. The probability of failure for lifetimes less than the 20-year target is also easily estimated and is plotted in Figure 3.

The number of components likely to fail is assessed by first setting the material property to a specified percentage-survival level. FAROW is then used to calculate the probability that the survival level will be achieved at several target lifetimes based on all the remaining uncertain quantities. The results of this analysis for several survival levels and a range of target lifetimes from 2 to 20 years are shown in Figure 4. While it is practically a sure thing that this component will exceed 50% or 60% survival rates even after 20 years, it has only a slim chance of achieving a 99.9% survival percentage at 20 years. The chances of achieving survival percentages between these extremes can be seen for this example in Figure 4.

Figure 4 plots cumulative probabilities in both directions: probability of some percentage or fewer components failing and probabilities of failure in less than some number of years. The probabilities can be broken down in various ways to provide a better insight into the nature of the risk by taking differences in both directions. Figure 5 shows the probability that the number of failures will be in a certain range of percentages of the entire fleet at certain of ranges of lifetimes. The figure is rendered less cluttered by only plotting the operating year ranges of less than 2, 8 to 10, and 18 to 20.

It is interesting that the probability of less than one percent failure rate is high in the first two years but goes to zero as the years advance. This shows the effect of common sources of uncertainty. Over time, two possibilities for fleet response emerge; either none of the components fail or small but significant numbers will fail. There is only a remote chance that there will be only rare failures (between 0 and 1%) at 20 years.

Expected Costs

There are an unlimited number of options for how costs can be related to component failures. The simplest is a constant cost model where the total cost is the number of failures times the cost of fleet replacement, including lost revenue from down time. This model may not reflect some important nonlinearities, such as an increase in per-unit lost revenue due to reduced availability when there too many failures to effect timely repairs. However, the constant cost model will be applied here for simplicity.

The expected cost is calculated by multiplying the cost of fleet replacement by the percentage of components failing (approximated by the center of the range) and multiplying by the probability of occurrence. The cost of replacing every component in the fleet is set at 100 for this example. The result of applying this cost model to Figure 5 is shown in Figure 6. Notice that because the cost increases with the number of components failing, the later years of operation show a marked increase in the expected cost of replacements, even though the probabilities were lower.

Cumulative costs in each percentage range are obtained by summing over all earlier year ranges and are shown in Figure 7. This example illustrates a case for which about half the expected cumulative costs are due to the risk that small numbers of components (less than 10% of the total installed) will have failed in 20 years of operation. It is perhaps more interesting that *half* the financial risk at 20 years comes from the chance that *more than 10%* of the components will fail. This outcome may not have been apparent because there is only a 3% probability that any particular component will fail in the 20-year target lifetime.

The total expected (average) cumulative costs are shown in Figure 8. Notice the similarity with Figure 3. With the constant cost model, the results are the same as if the probability of individual component failure is multiplied by the cost of replacement (i.e., the expected cost is 3% times the replacement cost of 100, or 3). The calculation of expected cost is the same whether the sources of the uncertainty are broken out or not. They are only the same, however, if a constant cost model is used. Costs that depend on the number of failures require breaking down the expected failures into ranges as was done here and would yield different results.

Incremental costs in each 2-year interval are calculated either by taking the difference in cumulative costs between consecutive 2-year increments or by summing over all the possible ranges of percentage failures in Figure 6. The present value of these incremental costs is shown in Figure 9 for discount rates of zero, 5% and 9%. The present values for 20 years of operation come out to be 3.1, 1.7 and 0.8, respectively. Notice that because the incremental costs initially increase and then level out, the present values are dominated by the costs in the later years, unless the discount rate is high.

Variability in Costs

Average costs, however, do not tell the entire story when it comes to projections dealing with uncertainty. The up-side (best case) and down-side (worst case) risks should also be known. The information in Figure 4 is used to estimate the cost associated with different levels of probability of occurrence. By slicing across

at a selected probability level, interpolating to find the percentage of components failing at that level in each year, and multiplying by the cost of replacing that percentage of the components, an estimate of the cumulative costs associated with the probability of occurrence is obtained for each year.

Figure 8 contains cost estimates at the 50th and 90th percentile levels plotted along with the expected, or average, cost. The 50th percentile reflects the level at which half of the possible outcomes lie on either side. Notice that the average cost is much higher than the 50th percentile due to the highly skewed distribution of possible outcomes, including both substantial probabilities of zero cost (no failures) and small probabilities of very high costs. The up-side in this example is zero cost. The down-side risk is illustrated by the cost above the 90th percentile, which indicates a 10% probability that the cost in twenty years will be greater than 8, i.e., about three times the expected cost.

EXAMPLE WITH HIGHER UNCERTAINTY

A second example in which some of the uncertainties are higher may be useful in contrasting some of the results of the first example. Here we take another example from the FAROW user's manual (Veers, et al., 1994). Details are again found in the user's manual with the inputs summarized as before by plotting the relative importance of the uncertain input groups. This example is intended to represent a case, perhaps earlier in the development process, in which there is a substantial uncertainty as to the stress response of the structure due to dynamic loading. The uncertainty in material properties and wind speed is about the same as in the first example, but the stress response uncertainty, as shown in Figure 10, is increased to the point where it accounts for 85% of the total. This condition is likely when there has been insufficient prototype testing to validate the predicted stress levels and stress concentration factors. The median lifetime for this example is 700 years and the probability of failure in less than ten years is about 11%.

The probabilistic calculations are summarized in Figure 11 similar to the first example results in Figure 4. Notice, however, that two of the scales in Figure 11 are different than Figure 4: "Years of Operation" runs

from 1 to 10 years, and “Percentage Survival” now runs all the way from 1% to 99%. The graph indicates that while the probability of everything surviving is higher than before, there is a small but nontrivial chance that nothing will survive, even in the first year. The higher uncertainty in loading and response is responsible for the increased spread in possible outcomes.

Figure 12 shows the down-side risk by plotting the 90th percentile cost along with the expected (average) cost over time. The 50th percentile cost does not even appear on this plot; it is nearly zero over the entire ten year range. The 90th percentile cost, while starting below the expected cost, inflates rapidly until at 10 years it is almost half the total fleet replacement cost and more than four times the expected cost. In this example, the cost break down clearly indicates the up-side and down-side risks. In the best case, nothing breaks, which may reflect a situation where too much conservatism has been built into the design at an unacceptably high cost. In the worst case large numbers of components fail, again at an unacceptably high cost. This contrasts the first example where in the worst case small but significant numbers of components will fail and you can expect the rare failure even in the best case.

SUMMARY AND CONCLUSIONS

The expected (or average) cost of failed components over the operating life of a fleet of wind turbines can be calculated in either of two ways. The easiest is to estimate the probability of individual component failure and multiply by the total replacement cost. A second approach is to separate (as well as possible) the sources of uncertainty into common and independent sources. The probability of different numbers of components failing is then estimated and the expected cost of replacement for each percentage range is calculated. These two approaches result in the same expected replacement cost if the individual replacement costs are independent of the number of components failing. The incremental costs in each year of operation estimated by either method are used to calculate the present value of replacement costs. Average projected costs, however, are usually insufficient when there is significant uncertainty involved.

Costs associated with various percentile levels of probability of occurrence are also calculated here to show the up-side and down-side risks. Two separate examples illustrate the difference in up-side and down-side risks that result from different levels of uncertainty in the inputs.

There are three main reasons for dividing up the probabilities into the number of components expected to fail in each year:

1. Variable Cost Model: If costs are not simply equal to the number of failures times the individual cost of replacement, then variable costs will need to be assigned based on how many components have failed.
2. Understanding the Nature of the Risk: Since probability of individual component failure is not equal to the percentage of components projected to fail, it is useful to break down the percentages of components failing and project the evolution through time of failure probability and cost to better understand the financial risk.
3. Describing the Variability in the Risk: Up-side and especially down-side risks can be estimated to illustrate the range of possible cost outcomes. Costs associated with various probability of occurrence levels are calculated.

REFERENCES

- ASCE, 1982, "Fatigue Reliability: Introduction," by The Committee on Fatigue and Fracture Reliability of the Committee on Structural Safety and Reliability of the Structural Division, *Journal of the Structural Division*, American Society of Civil Engineers, Vol. 108, No. ST1.
- Jackson, K., 1992, "Deriving Fatigue Design Loads from Field Test Data," *Proc. WindPower '92*, American Wind Energy Association, pp. 313-320.
- Spera, D. A., 1994, "Fatigue Design of Wind Turbines," *Wind Turbine Technology*, D. A. Spera, Ed., American Society of Mechanical Engineers, New York, pp. 547-588.

- Sutherland, H. J., P. S. Veers, and T. D. Ashwill, 1994 "Fatigue Life Prediction for Wind Turbines: A Case Study in Loading Spectra and Parameter Sensitivity," *Case Studies for Fatigue Education, ASTM STP 1250*, R. I. Stephens, Ed., American Society for Testing and Materials, Philadelphia, 1994, pp. 174-207.
- Sutherland, H. J., 1989, *Analytical Framework for the LIFE2 Computer Code*, SAND89-1397, Sandia National Laboratories, Albuquerque, New Mexico.
- Van Den Avyle, J. A., and H. J. Sutherland, 1989, "Fatigue Characterization of a VAWT Blade Material," *Eighth ASME Wind Energy Symposium*, SED Vol. 7, Berg and Klimas, eds., ASME, New York, pp. 125-130.
- Veers, P. S., S. R. Winterstein, C. H. Lange and T. A. Wilson, 1994, *User's Manual for FAROW: Fatigue and Reliability of Wind Turbine Components*, SAND94-2460, Sandia National Laboratories, Albuquerque, New Mexico.
- Veers, P. S., 1989, "Simplified Fatigue Damage and Crack Growth Calculations for Wind Turbines," *Eighth ASME Wind Energy Symposium*, P. C. Klimas and D. E. Berg, Eds., SED-Vol. 9, pp. 133-140.
- Veers, P. S., "Blade Fatigue Life Assessment with Application to VAWTs," *Journal of Solar Energy Engineering, Trans. of the ASME*, Vol. 104, No. 2, May 1982, pp. 106-111.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Category	Input
Material Properties	<ul style="list-style-type: none"> • S-N Coefficient (C) • S-N Exponent (b) $N = CS^{-b}$ <p>where N is the fatigue life and S is the cyclic stress level</p> <ul style="list-style-type: none"> • Ultimate Strength
Stress Response	<ul style="list-style-type: none"> • RMS vs Wind Speed Coefficient (σ_C) • RMS vs Wind Speed Power Law Exponent (p) $\text{RMS} = \sigma_C (V / V_C)^p$ <p>where V is wind speed and V_C is a characteristic wind speed</p> <ul style="list-style-type: none"> • Mean Stress • Cyclic-Stress-Amplitude Distribution Shape Factor • Stress Concentration Factor • Cyclic Frequency (Cycle Rate)
Wind Speed	<ul style="list-style-type: none"> • Annual Average • Weibull Shape Factor • Cut-out Wind Speed
Other	<ul style="list-style-type: none"> • Availability • Miner's Rule Summation at Failure (Traditionally 1.0)

Table 1: Inputs to a Wind Turbine Component Fatigue Life Analysis, using FAROW, that may be Uncertain

Figure 1: Typical S-N test results for identical specimens (Van Den Avyle and Sutherland, 1989).

Figure 2: Relative importance of the three sources of uncertainty for the example.

Figure 3: The probability of component failure grows with time as the fatigue damage accumulates.

Figure 4: Probability of achieving different component survival percentages from 2 to 20 years.

Figure 5: Probability of the percentage of components failing being in specific ranges during specific time intervals.

Figure 6: Cost in three two-year intervals divided into ranges of percentage of components failing.

Figure 7: Cumulative costs at 2, 10, and 20 years in each range of percentage of components failing.

Figure 8: Cumulative cost of replacement from expected (average) costs and at two percentile levels of probability of occurrence.

Figure 9: Present value of incremental expected costs using different discount rates.

Figure 10: Relative importance of sources of uncertainty for the higher uncertainty example.

Figure 11: Probability of achieving different component survival percentages from 1 to 10 years in the higher uncertainty example.

Figure 12: Cumulative cost of replacement from expected (average) costs and 90th percentile probability of occurrence for the higher uncertainty example.

Fig. 1

Fig. 2

Fig. 3

Fig. 4

Fig. 5

Fig. 6.

Fig. 7

Fig. 8

Fig. 9

Fig. 10

Fig. 11

Fig. 12