

Variance Estimates of Wind Plant Capacity Credit

Michael R. Milligan

*Prepared for
AWEA Windpower '96
Denver, Colorado
June 23- 27, 1996*



National Renewable Energy Laboratory
1617 Cole Boulevard
Golden, Colorado 80401-3393
A national laboratory of the U.S. Department of Energy
Managed by Midwest Research Institute
for the U.S. Department of Energy
under contract No. DE-AC36-83CH10093

Work performed under task number WE617010

June 1996

NOTICE

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.

Available to DOE and DOE contractors from:
Office of Scientific and Technical Information (OSTI)
P.O. Box 62
Oak Ridge, TN 37831
Prices available by calling (423) 576-8401

Available to the public from:
National Technical Information Service (NTIS)
U.S. Department of Commerce
5285 Port Royal Road
Springfield, VA 22161
(703) 487-4650



Printed on paper containing at least 50% wastepaper, including 20% postconsumer waste

VARIANCE ESTIMATES OF WIND PLANT CAPACITY CREDIT

Michael R. Milligan
National Renewable Energy Laboratory
1617 Cole Boulevard
Golden, Colorado 80401
U.S.A.

ABSTRACT

As the utility industry adapts to meet the changing regulatory and business climate, it is becoming increasingly important for utilities to identify and quantify the risks in various aspects of doing business. To reduce the risk of depending too heavily on one specific type of generation or fuel, generation expansion planning techniques are incorporating methods of portfolio diversification theory. Financial option theory is also used to evaluate the relative costs of building now or building later. Applying these theories to utility planning helps utilities assess risks in the emerging competitive environment.

Risk is typically measured as a variance. For example, the risk associated with an investment can be characterized by the rate-of-return variance. Many studies that calculate the capacity credit of a wind plant do not calculate its variance, and therefore ignore risk. A capacity credit that is calculated in this way can be far different than the long-term average value. This problem is compounded by the usual method of calculating capacity credit, which depends very heavily on the level of wind generation during the system peak hours. A small change in wind power during the peak can have a dramatic effect on the capacity credit.

This problem is further compounded by the limited availability of multi-year wind data sets that can be used in utility production cost modeling. For example, a study that uses a single year of data and finds a 30% capacity credit may be based on a wind generation pattern that is not at all typical. Although the preferred approach would be to use many years of wind data to obtain a range of capacity credit estimates, this is not always possible.

This paper describes a technique that can help generation planners evaluate the variance of the capacity credit for wind power plants when there is limited wind data, and also shows some results of these calculations.

INTRODUCTION

One of the most frequently asked questions about wind power plants is whether such a plant has a capacity value, and if so, how much. Unfortunately, the answer to this question is rarely obvious, although it is often framed with the statement that capacity value depends heavily on the wind regime, utility load characteristics, and the utility's existing generation mix. There are also different methods that can be used to calculate capacity credit. Utility capacity expansion models and reliability models differ in their assumptions, algorithms, and their ability to properly account for a resource that is highly variable on both ends of the time scale. Utilities generally use 2 broad types of models: production cost and reliability. Generation expansion models are often built on the framework of production cost and reliability models. In some cases, the same model can be used to produce both reliability and power production cost outputs.

One of the most critical shortcomings of standard techniques used to measure wind plant capacity value is due to the variability of the resource and the lack of adequate wind data. This makes it difficult for the models to adequately measure capacity credit, so that capacity credit results may have little meaning. Because of the temporal interactions between load, wind power, and conventional generating capacity, wind plant capacity credit

measures are often little more than random draws from a probability distribution whose characteristics are largely unknown. To properly account for the exceedingly large number of potential interactions, some form of Monte Carlo simulation appears to be necessary. An excellent discussion of this technique in the context of chronological production cost models can be found in Marnay and Strauss (1990). However, in spite of the falling cost of computing resources, many production cost models have limited, if any, Monte Carlo capability. This paper uses a Monte Carlo technique that is carried out exogenously from the production cost/reliability model, which is then executed for many scenarios. Although this is not as computationally efficient as it would be if the capability were part of the production/reliability model, it can be applied to a wide variety of models that possess the capability of running a large number of scenarios.

It is also possible, if not likely, that long-term measures of capacity credit will differ from short-term measures. The focus of this paper is on long-term measures that would be appropriate for utility planners or investors who are evaluating a potential future wind plant. Short-term capacity credit, although outside the scope of this paper, will be mentioned again briefly below.

The usefulness of the concept of "wind plant capacity credit" has recently been questioned (Utility Wind-Modeling Planning Meeting, 1996). Citing the evolving deregulation of the utility industry, critics have argued that utility planning and capacity expansion will be influenced only by the market, reducing or eliminating the need for traditional capacity analysis. Under this scenario, capacity credit is determined by the pool or independent system operator (ISO) and not by traditional utility analysis. However, the final outcome of the deregulation process is not anything if not unclear. Although a number of states have begun moving toward a competitive market for electric utilities, the incentive to deregulate appears to be somewhat dependent on the price of electricity. States in the northeast and California generally pay the most for electricity, and that is where much of the deregulation effort has progressed significantly. The extent of federal regulatory involvement is also unclear. This could result in a patchwork of competitive and quasi-competitive markets for electricity.

If we assert that competition will indeed be pervasive and consistent, who plans for additional generating capability? It is the investors who are driven by the market. Investment in new generation would be driven by the expected rate of return that can be earned by the productive resource. To evaluate alternative investments, the potential investor must carry out calculations that allow the comparison of returns on these possible investments. This would most certainly include an estimate of the capacity payments that could be earned by a wind (or any other) power plant. If the investor is a generating company, the calculations that are carried out could conceivably be the same as those outlined in this paper.

METHODS USED TO MEASURE CAPACITY CREDIT

The focus of this paper is to examine capacity credit in the context of generation planning or investment. However, it is important to link the concepts of planning capacity credit and operational capacity credit. Planning capacity credit is the value given to a generating plant over a long time horizon, and is typically in the context of utility generation planning. Operational capacity credit is the capacity value that could be specified in a transaction between utilities. Utility A might agree to provide Utility B with 50 MW according to a pre-arranged schedule during a particular day or week. If this capacity can be provided by a wind plant, then the wind plant is said to have an operational capacity credit of 50 MW during the appropriate period. This section provides a short discussion of both types of capacity credit.

The standard techniques that are used to evaluate the reliability of power systems and how these techniques are used to measure planning capacity credit are based on Billinton & Allan (1984). Most methods of assessing the capacity credit of a wind plant are based on a reliability measure called loss of load expectation (LOLE). Most production cost and generation expansion models calculate the LOLE or a related measure, such as loss of load hours or expected unserved energy. Although these measures are not equivalent, they are measures that capture

the possibility that the generating system is not adequate to meet the system load. Of course the goal of the utility is to keep this probability as small as possible, given the trade-off between cost-minimization and reliability. A standard rule-of-thumb is to maintain an expected loss-of-load expectation of 1 day in 10 years.

There are other ways in which a utility can gauge its reliability. Another approach is to maintain a reserve capacity margin that exceeds peak load by a given percentage. Although there is no direct formula for converting between reserve margin and LOLE, higher reserve margins correspond to a lower LOLE and hence a more reliable system.

Using the concepts and techniques from reliability theory (Billinton and Allan, 1984), we want to provide a measure of generating plant capacity credit that can be applied to a wide variety of generators. Although no generator has a perfect reliability index, we can use such a concept as a benchmark to measure real generators. For example, a 500-MW generator that is perfectly reliable has an effective load carrying capability (ELCC) of 500 MW. If we introduce a 500-MW generator with a reliability factor of .85, or equivalently, a forced outage rate of .15, the ELCC of this generator might be 390 MW. In general, the ELCC value cannot be calculated by multiplying the reliability factor by the rated plant output — the ELCC must be calculated by considering hourly loads and hourly generating capabilities. This procedure can be carried out with an appropriate production-cost or reliability model.

To find the ELCC of a new generator, one must evaluate the reliability curve at various load levels prior to adding the new generator to the system. This can be done by running the reliability model, altering the load, and plotting the resulting points in a graph such as that in Figure 1, below. The graph shows the increasing risk of not meeting load, as measured with LOLE, that results from load increases. In the figure, the system load-carrying capability is just under 1,100 MW, assuming a risk level of 1 day in 10 years. The utility that finds itself above its preferred level of risk would add generation to its system. The new generator would shift the reliability curve to the right. This is depicted in Figure 2. The level of load increase that can be sustained at the same reliability level is the distance between the 2 risk curves, evaluated at the preferred risk level. Later discussion in this paper will use this method to determine the ELCC of a wind plant.

The determination of short-term operational capacity credit is a different process. If a wind-plant operator contracts with a utility to provide capacity on a given schedule for a given day, it is in the best interest of the wind-plant operator to possess a consistently accurate forecast of the wind, and hence windpower availability, during the day in question. In the unlikely but optimal case, the wind forecast is known with absolute certainty.




FIGURE 1. RELIABILITY CURVE FOR A FICTITIOUS UTILITY




FIGURE 2. RELIABILITY CURVES FOR CALCULATING ELCC OF WIND PLANT

In that case, the wind-plant operator would contract for the full available capacity. In this case, the capacity value of the wind plant is equal to the capacity level that can be sold on a firm basis for the day (or any appropriate time interval) in question. During windy periods of the year, this capacity level is likely to be relatively high, whereas in the calm seasons this capacity level will be lower. The operational capacity credit can therefore vary throughout the year, and can be summarized by a suitable annual average, along with a variance measure. Of course, the forecast will contain an error component. The contract negotiators must quantify the relative risks of aiming too high or too low. A more detailed discussion can be found in Milligan, Miller, and Chapman (1995). However, in the "long run" we would expect that the average of the operational capacity credits would approach the long-term capacity credit, as measured later in this paper.

WIND PLANTS AND RELIABILITY

Adequately representing wind power systems in hourly reliability and production cost modeling can present a challenge, particularly if the model uses the load duration curve (LDC) approach. As computing platforms have become more powerful over the past few years, there has been additional interest in chronological models. However, much of the early work of calculating wind plant capacity credit was done with LDC models. In the LDC framework, loads are grouped into subperiods that consist of some reasonable partitioning of the hours in a week or month. The loads are sorted, and used to calculate a probability density function that is used to find the economic dispatch or reliability values of interest. This process eliminates a significant computational burden, but does so by sacrificing the chronological nature of the load data. Because the correlation between wind power and customer load is important to capture in the modeling, analysts have typically subtracted the hourly available wind power from the load. The result of this set of calculations is the remaining load, which is then met with the usual rules of unit commitment and economic dispatch (although the latter is not typically found in reliability models).

A similar technique for calculating net equivalent load can be used with chronological models. The justification is that a least-cost dispatch strategy will always take an inexpensive variable-cost resource, such as wind, before more expensive options. After wind power is accounted for, the conventional generating resources can be called upon to meet the remaining load. The chronological model overcomes the time-scale limitation of the LDC model. However, treating wind power as a deterministic reduction in load poses the same problem for chronological models as with LDC models.

Capacity credit results depend heavily on what happens during the utility's peak hour or several peak hours. Wind speed varies significantly from year to year and from hour to hour. Capacity credit estimates that are based on a single year of data and modeled without taking this variation into account should be suspect. Some analysts have corrected for this problem (Percival and Harper, 1982), whereas others did not (Bernow, Biewald, Hall, and Singh, 1994). A recent paper by Billinton, Chen, and Ghajar (1996) takes an approach that is similar to this paper. Ignoring this problem can be perilous, and can result in significantly over- or under-estimating capacity credit.

As an example of the wide potential variation in year-to-year wind energy capture, I have done a brief analysis of a 14-year data set from a regional air quality monitoring program (RAMP) site in North Dakota. It is important to note that this site would not be judged as suitable for a wind power plant, because of its low average wind speed and other factors. However, the data series is composed of many years, and until more multi-year data sets are publicly available from potential or actual wind plant sites, it is useful to look at this series.

To illustrate the possible variation in annual energy capture, this 14-year data set was used to calculate annual energy for a fictitious wind plant. The results are presented in Figure 3.

As the figure indicates, there is wide variation in annual energy capture. In 1983, for example, annual energy produced from this site would be less than 60% of that produced in 1988. This clearly points out the fallacy of using a single year of wind data for meaningful analysis. When several years of data is not available, what then? That question is addressed in the remainder of this paper.




FIGURE 3 - ANNUAL SIMULATED ENERGY PRODUCTION, RAMP DATA SITE

MODELING APPROACH

To provide a plausible analysis of wind plant reliability and ELCC, I apply a Markov (Hillier and Lieberman, 1974) wind-speed simulation tool to a single year of wind data. The wind data is from the U.S. Department of Energy Candidate Wind Site program (Sandusky et. al. 1983). The site chosen for this work is from Romero Overlook, scaled to match the Altamont Pass site. For the utility's peak month, a state transition matrix is calculated. Then multiple realizations of the data are calculated by repeatedly sampling from the state transition matrix. This technique preserves some of the time-scale properties of the wind speed data and also provides an estimate of the variation that could be expected from a wind site. This method suffers from an obvious limitation—only a single year of wind data is used to calculate the state transition matrices. Including additional wind data, if available, would increase the accuracy of these calculations.

This analysis focuses on the month of the utility system peak, although the method could be used on any appropriate time-frame. Some utility control areas, pools, or reliability regions estimate generating plant capability on a monthly basis, so the choice of time-frame is consistent with those approaches. Once the multiple wind speed series have been simulated, I can calculate the hourly wind power output for the month from a hypothetical wind plant from each realization. The hypothetical wind plant is then applied to a reliability model, which is executed for each wind plant realization, and the results are combined. From this process I obtain the ELCC of each wind plant realization, which can then be summarized for further analysis.

The tools used in this study include Wind Power Simulator, described in some earlier work (Milligan & Miller, 1993), and the Elfin production cost model (Elfin is a product of the Environmental Defense Fund). For these cases the economic dispatch logic of Elfin was overridden so that I could focus on installed reliability using all generating resources. Additional software tools were used to simulate the multiple wind speed realizations and

summarize the multiple Elfin outputs. The utility data is from an actual utility, modified appropriately for this study. Figure 4 below provides a graphical depiction of the process.

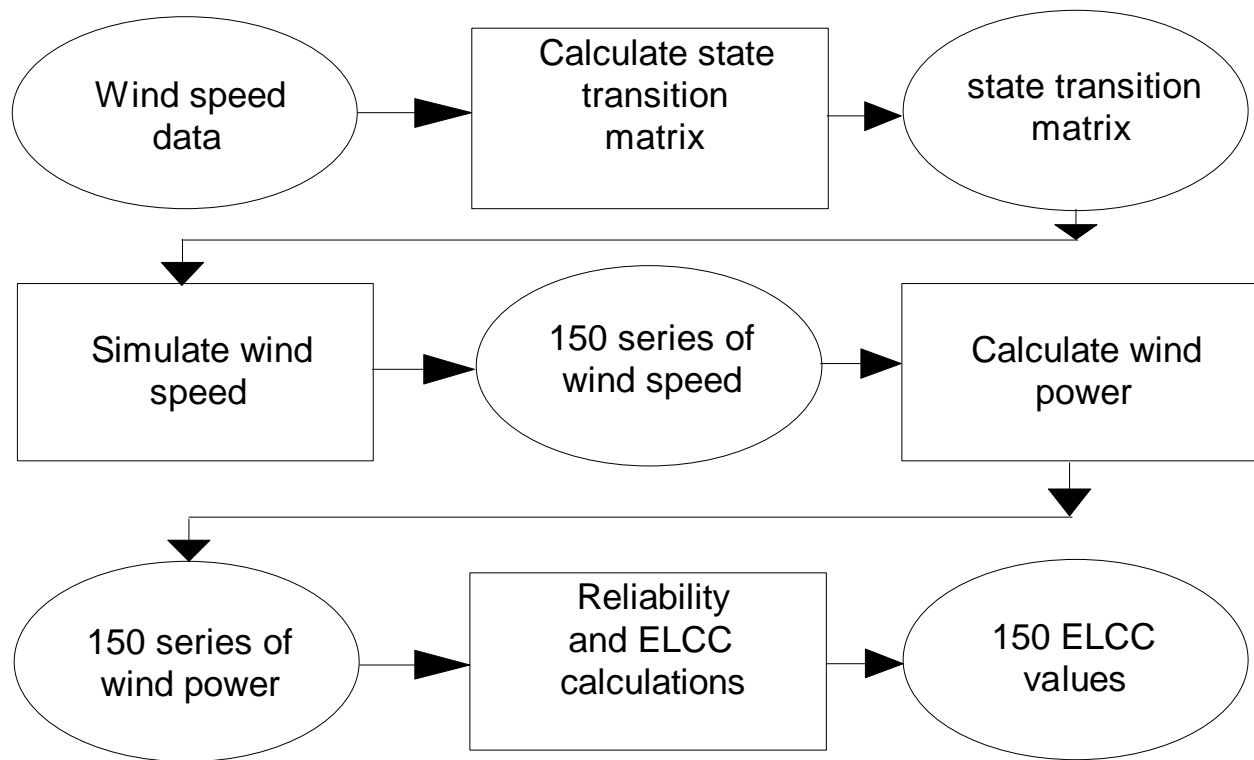


FIGURE 4. FLOWCHART OF THE MODELING PROCESS

Modeling wind power plants in production cost or reliability models requires the modeler to make many decisions about how the plant should be characterized for the model itself. One of the most important decisions is whether the wind plant capacity should be counted as "firm" or "non-firm." A generating unit that is modeled as a non-firm unit can't contribute to the utility's commitment target. By definition, non-firm resources do nothing to improve the reliability calculation, even though they may marginally improve actual reliability. If a unit is treated as non-firm its output is likely to be curtailed on very short notice. Although this situation does not arise often in practice, it implies that another unit must carry spinning reserve to cover the potential outage. Assuming a partially accurate wind forecast, wind plants should not be modeled as non-firm. The purpose of this study is to determine the capacity value and its variation. The designation of a firm vs. a non-firm resource tells the model how a particular resource should be treated for the reliability calculation. For this study, the wind plant was modeled as a firm resource, indicating that its full hourly capacity should be counted in the reliability calculation. Because I calculate a full range of possible outcomes with multiple wind data sets, this procedure allows me to capture such measurements as average capacity on peak or variation of capacity on peak. It also allows me to perform the capacity credit calculation based on the many cases that have been run. For a more detailed discussion of firm and non-firm treatments of wind plants and the relationship to operational capacity credit and wind forecasting, see Milligan, Miller, and Chapman (1995).

RESULTS

The peak-month state transition matrix appears in Figure 5. The three-dimensional graph shows the probability density function for each pair of successive wind velocities. The location of the peaks on the density graph shows a high level of auto-correlation in the wind-speed series.




FIGURE 5. PEAK MONTH STATE TRANSITION MATRIX

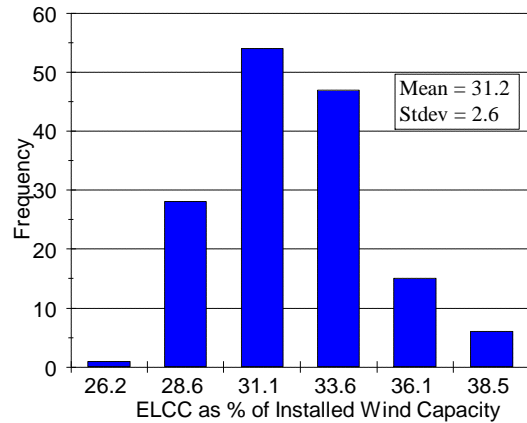


FIGURE 6. CAPACITY CREDIT DISTRIBUTION FOR 150 RUNS

Figure 6 captures the results of the multiple simulations. For each of the 150 runs, ELCC was calculated as a percentage of installed wind plant capacity. The data from these calculations is what the figure summarizes. The intervals were chosen to approximate the standard deviation of the ELCC value. As the figure indicates, most of the ELCC values are within 1 standard deviation of the mean value of 31.2%, with approximately 20 values outside that range. The ELCC values for this wind plant range from a minimum of 25.6% to 37.7%, a difference of 40% of the mean ELCC value.

I have developed reliability curves for several cases. The first case is the mean case, and it shows the simulation case whose ELCC value most closely approximates the mean ELCC value. The second and third cases,

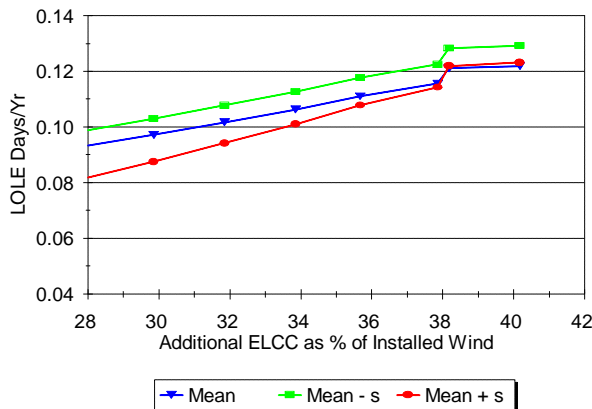


FIGURE 7. RELIABILITY CURVES WITHIN ONE STANDARD DEVIATION OF MEAN

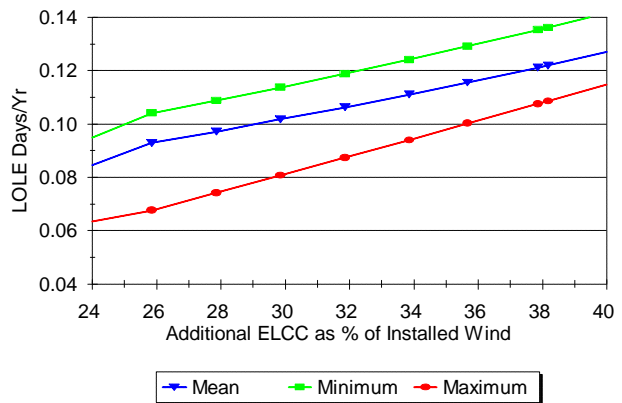


FIGURE 8. RELIABILITY CURVES FOR MEAN, MINIMUM, MAXIMUM ELCC CASES

respectively, are those whose ELCC values most closely match the mean plus or minus the sample standard deviation. These reliability curves are combined in Figure 7. The distance between these cases at the risk level of 1 day/10 years is the variation in ELCC, which appears in Figure 6.

It is also useful to look at the reliability curves for other extreme cases. Figure 8 shows the average ELCC case along with both the maximum and minimum cases. It should be apparent from this graph that decision-makers and analysts should be extremely cautious about representing the capacity value of a wind plant as a single, fixed value.

It is also of interest to examine the convergence of the ELCC calculations as the number of scenarios (or iterations) increases. Adopting a method similar to that used by Billinton, Chen, and Ghajar (1996), I calculate the cumulative average and standard deviation of ELCC for each iteration. The ratio of the standard deviation to the cumulative mean is calculated. This ratio is called the convergence factor, and is graphed in Figure 9.

It is apparent from Figure 9 that a near steady-state value is reached in less than 100 simulations. Although the cost in computer run-time for this number of iterations is somewhat expensive, there is clearly an important benefit derived from running even a smaller number of iterations, as one can obtain a better idea of the possible variation in capacity credit. Figure 10 shows the percentage change in the convergence factor as a function of the number of iterations.

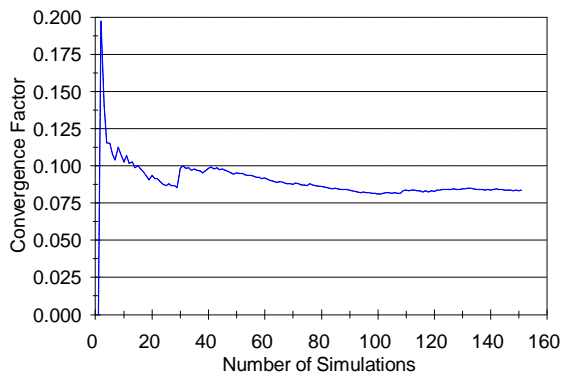


FIGURE 9. CONVERGENCE OF THE SIMULATIONS

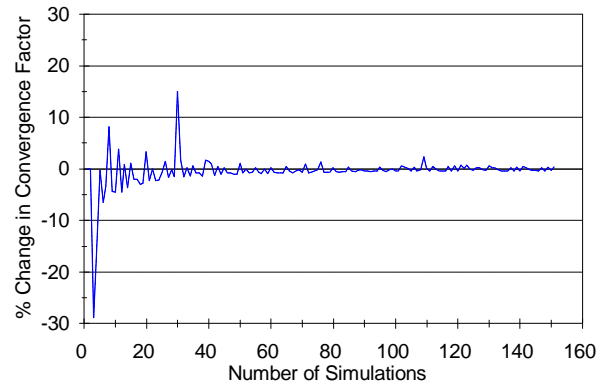


FIGURE 10. PERCENT CHANGE IN CONVERGENCE FACTOR

These results are similar to those in Billinton et al., but these results converge much faster. The Billinton wind-speed model is based on an autoregressive moving average (ARMA) model, described as an ARMA(3,2) process. A simpler form of the ARMA model, the AR(1) model, is a special case of the Markov model. It is possible that the ARMA model more accurately characterizes the time series data, which might explain the large number of iterations required by the ARMA model to converge—about 6,000.

CONCLUSIONS

This paper illustrates a Monte Carlo simulation technique that can be used to estimate the capacity value of a wind plant when there is limited wind data available. It is apparent that a single-year sample of wind data is inadequate to properly evaluate capacity credit unless some form of repeated sampling is used. The Markov wind-speed model used in this paper provides a relatively simple procedure to examine potential variation in capacity credit. The technique can also be applied to any other model output of interest, such as production cost

or fuel usage. However, the use of multi-year data sets to calculate the state transition matrix would likely introduce more variation into the capacity credit measure than I have calculated here.

The case study results indicate that wind power plants do indeed have capacity value, and that this value is subject to a reasonably wide variation. Of course the specific results are influenced by the wind regime, utility loads, other conventional generation, and other characteristics of the utility.

REFERENCES

- Bernow, S., B. Biewald, J. Hall, D. Singh (1994), "Modelling Renewable Electric Resources: A Case Study of Wind." *Tellus* No. 91-187. Boston, Massachusetts. Tellus Institute.
- Billinton, R. and R. Allan (1984). *Reliability Evaluation of Power Systems*. Plenum Press. New York.
- Billinton, R. H. Chen, and R. Ghajar (1996), *A Sequential Simulation Technique for Adequacy Evaluation of Generating Systems Including Wind Energy*. IEEE/PES Winter Meeting, January 21-25, 1996. Baltimore, Maryland.
- Hillier, F. and G. Lieberman (1974), *Operations Research*. Holden-Day. San Francisco.
- Jarass, L., L. Hoffmann, A. Jarass, G. Obermair (1981), *Wind Energy: An Assessment of the Technical and Economic Potential*. Springer-Verlag. New York.
- Marnay, C. and T. Strauss (1990). *Variance Reduction in Monte Carlo Chronological Production-cost Modeling*. California Public Utilities Commission.
- Milligan, M. (1996). *Alternative Wind Power Modeling Methods Using Chronological and Load Duration Curve Production Cost Models*. NREL/TP-441-8171. Golden, Colorado. National Renewable Energy Laboratory.
- Milligan, M. (1996). "Measuring Wind Plant Capacity Value." *WindStats Newsletter; Vol. 9, No. 1*; Knebel, Denmark. NREL/TP-441-20493. Golden, Colorado. National Renewable Energy Laboratory.
- Milligan, M. and J. Klarner (1995). *Wind Data Replacement Algorithms*. Seminar presented at the National Renewable Energy Laboratory, Golden, CO, 19 December.
- Milligan, M. and A. Miller (1993). "The Value of Windpower: An Investigation Using a Qualified Production Cost Model." *Windpower '93 Proceedings; July 12-16, 1993*; San Francisco, California. Washington, DC: American Wind Energy Association; pp. 74-81.
- Milligan, M., A. Miller, and F. Chapman (1995). "Estimating the Economic Value of Wind Forecasting to Utilities." *Windpower '95 Proceedings; March 27-31, 1995*; Washington, DC. Washington, DC: American Wind Energy Association.
- Percival, D., and J. Harper (1981), *Value Analysis of Wind Energy Systems to Electric Utilities*. SERI/TP-732-1064. Golden, Colorado: Solar Energy Research Institute.
- Sandusky, W., J. Buck, D. Renne, D. Hadley, O. Abbey, S. Bradymire, J. Gregory (1983), *Candidate Wind Turbine Generator Site Cumulative Meteorological Data Summary and Data for January 1982 Through September 1982*. PNL-4663. Richland, Washington. Pacific Northwest Laboratory.
- Utility Wind-Modeling Planning Meeting. (1996). National Renewable Energy Laboratory, Golden, CO, 22-23 February.