NREL/TP-440-21530

An Enumerated Probabilistic Simulation Technique and Case Study: Integrating Wind Power into Utility Production Cost Models

Michael R. Milligan
National Renewable Energy Laboratory

Mark S. Graham
Tri-State Generation and Transmission Association, Inc.

Presented at
IEEE Power Engineering Society
Summer Meeting
Denver, Colorado
July 29–August 1, 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

July 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

An Enumerated Probabilistic Simulation Technique and Case Study: Integrating Wind Power Into Utility Production Cost Models

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

Mark S. Graham
Tri-State Generation and Transmission Association, Incorporated
1100 W. 116th Avenue
P.O. Box 33695
Denver, Colorado 80233
U.S.A.
margra@tristategt.org

Abstract

Production cost, generation expansion, and reliability models are used extensively by utilities in the planning process. However, many of these models do not provide adequate means for representing the full range of potential variation in wind power plants. In order to properly account for expected variation in wind-generated electricity in these models, we describe an enumerated probabilistic approach that can be performed outside the production cost model, compare it with a reduced enumerated approach, and present some selected utility results. Our technique can be applied to any model, and can result in a considerable reduction in model runs. We use both a load duration curve model and a chronological model to measure wind plant capacity credit, and also present some other selected results.

Introduction

Representing wind power plants in utility production cost and reliability models poses a challenge to modelers because of the wide range of potential variability of the resource. As utilities evaluate wind power plants for possible future resource additions, it is important to accurately quantify both the capacity value and energy value of the wind plant. Renewable power plants, such as wind, may also contribute other benefits. Among these are fuel diversity and lack of emissions. In this paper we illustrate two related techniques that can be used to help capture some of this variability, and apply the technique to measuring capacity credit. The first method was introduced by Milligan (1996a), and is an enumerated probabilistic approach (EPA). The second method is based on a selective reduction in the multiple models runs, and is called the reduced enumerated probabilistic approach (REPA).

Capacity credit measures of wind power plants help utility planners and decision-makers evaluate this intermittent resource in the context of other types of power plants. The term "capacity credit" refers to that level of conventional generation that can be replaced with wind generation. To avoid complications of how to compare the many types of conventional generation, analysts will either compare wind to an ideal generation source or to a specific type of generator.

There are many possible techniques that can be used to calculate wind plant capacity credit. The choice of method is influenced by several factors, such as data availability and scope of the study. When capacity credit calculations are performed using a production-cost or reliability model, it becomes possible to do more precise calculations so that the wind power plant capacity credit is measured by a reliability index. The most common approach calculates the effective load carrying capability (ELCC) that is described by Billinton and Allan (1984), Garver (1966), and others.

Two of the most critical shortcomings of the standard techniques used to evaluate wind plant capacity value are 1) variability of the resource and 2) the lack of adequate wind data. Intermittency and the high variability of wind makes it difficult for models to adequately measure capacity credit, so 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 large number of potential interactions, some form of Monte Carlo simulation is necessary. An excellent discussion of this technique in the context of chronological production cost models can be found in Marnay and Strauss (1990). Milligan (1996a) illustrates a Monte Carlo method that is external to the load-duration curve production cost model. This approach allows for the creation of an enumerated set of wind power series, each of which can be run in the production cost or reliability model. This process is called the enumerated probabilistic approach (EPA) to differentiate it from the intrinsic Monte Carlo approaches that can sometimes be found in production cost and reliability models. However, in spite of the falling cost of computing resources, many production cost models have limited, if any, Monte Carlo capability. Other models may possess a modest Monte Carlo capability, but one that is not capable of providing reasonable sampling of wind-power output. This paper uses the EPA that is implemented outside the production cost/reliability model, which can then be executed for any number of scenarios. We then illustrate a variation of the EPA method that uses a reduced number of enumerated cases. This is called the Reduced EPA, or REPA, and is carried out in weighted and unweighted variations. The advantage of the REPA methods is that they are not as computationally demanding, although they will result in a modest loss of accuracy.

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 clear. 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 wind plant's contribution to system reliability and the capacity payments that could be earned by the 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.

This paper illustrates a technique that can be applied to any production-cost/reliability model that is capable of handling multiple wind power series, and extends earlier work by Milligan (1996a). We use two production-cost models: Elfin, a load-duration curve model, and P+, a chronological model. These models are described below. Using the Elfin model to establish a base case of 100 EPA simulations, we then use both Elfin and P+ to perform a set of REPA simulations that is less computationally intensive. The paper concludes with some selected results and comparison of the EPA and REPA results.

Methods Used to Measure Capacity Credit

One focus of this paper is to examine capacity credit in the contexts of generation planning or investment, and plausible variations of the wind resource. 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 used 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 the way these techniques are used to measure planning capacity credit are based on Billinton and 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 common practice is to maintain an expected loss-of-load expectation of 1 day in 10 years. Other reliability indices can be used in place of LOLE. One such measure is expected unserved energy (EUE). The advantage of this measure for our work is that both of our models produce a measure of EUE, allowing us to legitimately compare the outputs of both models. See Billinton and Allan (1984) for a discussion of reliability indices.

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 or EUE, 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 with various load levels and plotting the resulting points in a graph.

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 wind power availability, during the day in question. It would be optimal but unlikely that the wind speed is known with absolute certainty. The capacity value of the wind plant is 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). 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, Reliability, and Capacity Credit

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 for this technique 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 singular deterministic reduction in load poses the same problem for chronological models as with LDC models: system reliability measures with respect to the wind resource are not accurate.

Capacity credit results depend heavily on what happens during the utility's peak hour or several peak hours. Wind speed can vary 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, 1981), 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 Milligan (1996a), which we extend here. 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, we have done a brief analysis of a 13-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 13-year data set was used to calculate annual energy for a fictitious wind plant. The results are presented in Figure 1.

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 Milligan (1996a), and we expand on that here.

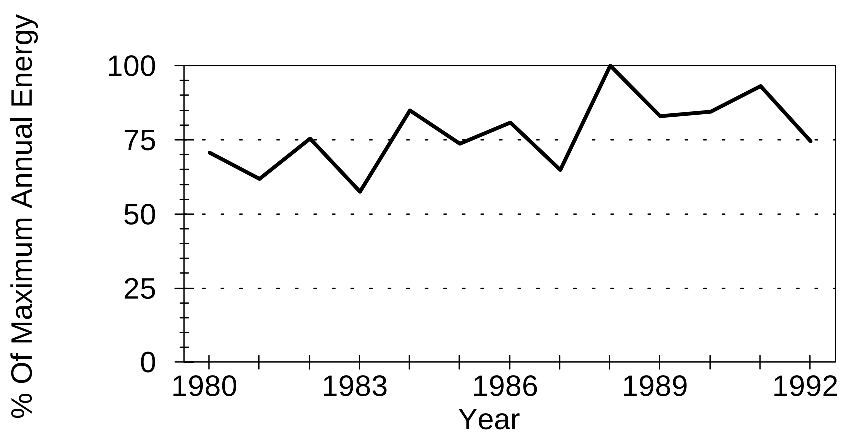

Figure 1. Annual simulated energy production, RAMP data

Modeling Approach

The utility data we used is from Tri-State Generation and Transmission Association, Inc. Tri-State is a non-profit generation and transmission cooperative, supplying wholesale electric power to 34 distribution cooperatives in Colorado, Wyoming, and western Nebraska. Resources include both Tri-State-owned and jointly-owned coal and oil-fired generation. Tri-State also purchases power from the Western Area Power Administration (Western) and Basin Electric Power Cooperative (Basin).

To provide a plausible analysis of wind plant reliability and ELCC, we apply a Markov (Hillier and Lieberman, 1974) wind-speed simulation tool to a single year of wind data. Other similar Markov applications can be found in Manwell, Deng, and McGowan (1994) and Deshmukh and Ramakumar (1982). The wind data is from the Nebraska Energy Office. We chose the Imperial, Nebraska site because of its proximity to Tri-State's service territory. For each 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 reasonably 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 ability to represent long-term data, or, in the limit, negate the need for a wind-speed simulation tool altogether.

This analysis focuses on October, 1995, a month in which there appears to be significant variability in the wind resource. To satisfy the requirement that partial weeks are not allowed, we ran each model for 6 weeks and obtained calendar summaries for October. Some weekly results are reported below. For this month Tri-State's peak load was 1,440 MW. To minimize differences between production models we reduced this load by 90 MW to account for a time-varying purchase from Basin. Net peak load was 1,350 MW. The maximum hydro purchase from Western was 400 MW, with 1,152 MW of base and intermediate generation and 120 MW of peaking capacity. We modeled a hypothetical 100 MW wind plant.

The wind-speed state transition matrix for October appears in Figure 2. This graph shows the probability of occurance of each wind-speed at time *t* as a function of velocity at time *t-1*. Our 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, we can calculate the hourly wind power output for the month from a hypothetical wind plant from each realization. We can then perform the analysis of either the full Monte Carlo or the weighted Monte Carlo, both of which we describe below.

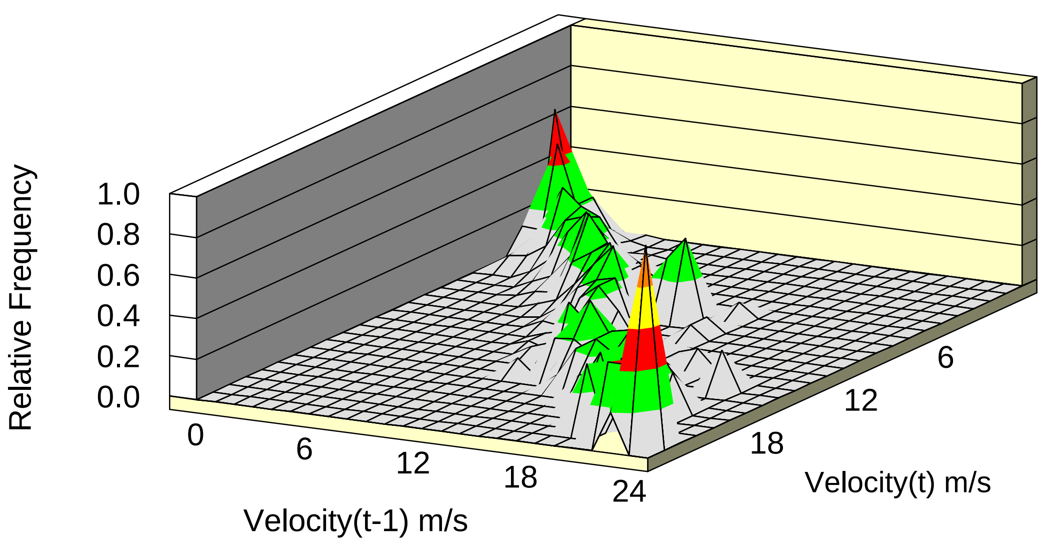

Figure 2. Wind speed state transition matrix for Imperial, Nebraska, October 1995

For the EPA runs, each hypothetical wind power series is input to the production-cost/reliability model, which is executed for each one. From this process we obtain the ELCC of each wind plant realization, which can then be summarized for further analysis. The Elfin model, described below, is used for this method. Figure 3 below provides a graphical depiction of the EPA process.

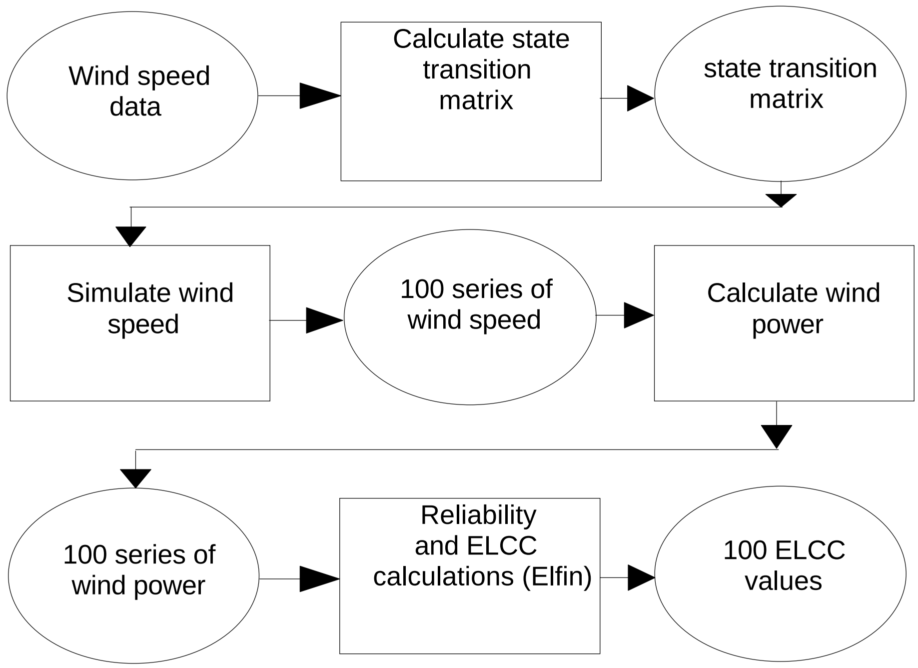

Figure 3. Enumerated Probabilistic Approach

The REPA method is an attempt to reduce the number of reliability model executions with a minimal reduction of accuracy. Our approach is to analyze the 100 wind-power series based on energy output during the utility system peak. The data are then grouped and weighted. Representative wind-power series are then selected from each of the groups. The reliability model is then run once for each group. For the weighted REPA the weights are applied to the output of interest. We performed this analysis with both the Elfin and P+ models. Figure 4 shows the weighted REPA approach.

Modeling Tools

The tools used in this study include Wind Power Simulator, described in an earlier work (Milligan and Miller, 1993), the Elfin production cost model (Elfin is a product of the Environmental Defense Fund), and the P+ model (P+ is a product of the P-Plus Corporation), both briefly described here. Additional software tools were used to simulate the multiple wind speed realizations and summarize the various model outputs.

The Elfin model uses the load-duration curve method for calculating production cost and reliability. Weeks can be divided into 14 subperiods that correspond to the utility's peak and off-peak periods. Elfin uses piecewise, linear LDC for each weekly subperiod. Elfin has several options for modeling thermal and hydro generation. Resources with critical time-profiles, such as hydro and wind, can be modeled as time-varying load-modifiers. For these resources Elfin applies the resource to the chronological load prior to constructing the LDC. Elfin performs unit commitment on a weekly basis, and provides the user with various options for modifying unit commitment for particular machine configurations.

The P+ model is a derivative of the PowerSym model, originally developed at the Tennessee Valley Authority. The model has been further developed and enhanced by the P-Plus Corporation. The P+ model can produce output for hours, days, months, and years. The model allows for several thermal and hydro generator types, and commitment and dispatch occur in sequence by type. Wind can be modeled as an hourly transaction or as an equivalent generator with a specified hourly must-run schedule. The model is chronological, and allows the user to specify unit ramp rates and minimum up- and down-times. See Milligan (1996b) for additional details.

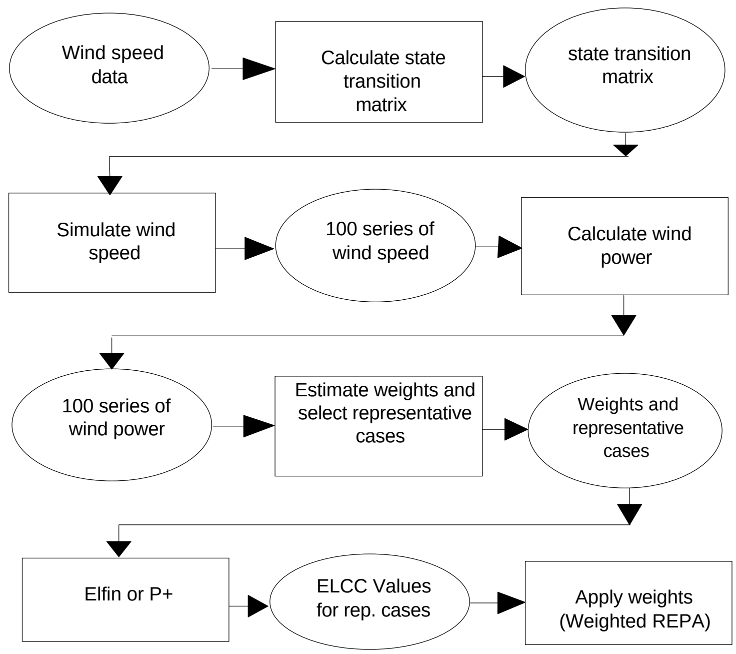

Figure 4. Reduced Enumerated Probabilistic Approach (REPA) and weighted REPA

Representation of Wind Power Plants in Elfin and P+

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 versus 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 we calculate a full range of possible outcomes with multiple wind data sets, this procedure allows us to capture such measurements as average capacity on peak or variation of capacity on peak. It also allows us 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).

We modeled wind power as a load modifier in Elfin and as a fixed hourly transaction in P+. This approach causes each model to treat the wind power plant in the same way. The hourly load is reduced by the level of wind generation in that hour, and conventional resources are committed and dispatched accordingly.

Weight Selection for the Weighted REPA

The process of grouping the various wind scenarios involves some judgement. Our intent is to select the data bins in such a way that the variation of the binned data closely represents the variation in the ungrouped data. In our judgement, grouping the data with bin sizes corresponding to the mean and standard deviation resulted in poor representation of the variation we found in the ungrouped cases. Faced with the inevitable tradeoff between execution time and accuracy, we did not want to use a large number of bins, since the saving in model runtime compared to the EPA method would not be significant. Figure 5 illustrates the distribution of the wind energy during the utility's peak period in October. To properly interpret each bin, note that the x-axis labels designate values up to and including each respective bin label.

Figure 6 shows the bins and resulting weights that we selected. This grouping retains the bimodal nature of the original distribution, while adequately representing the variation in the data.

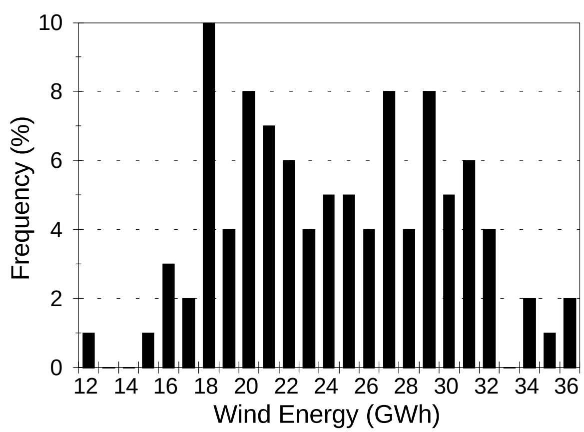

Figure 5. Wind energy distribution for multiple data sets for October

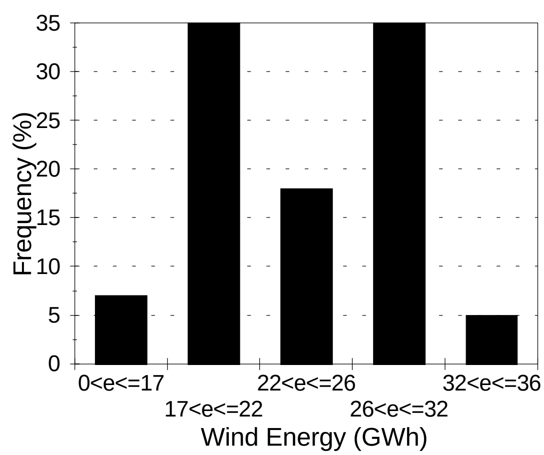

Figure 6. Wind energy distribution for multiple data sets, binned, with relative weighting factors

Capacity Credit Results

After making the adjustments to the load data as described above, both Elfin and P+ were executed to obtain base-case results with no wind generation. Table 1 shows the reliability outputs from each model. It is clear from the table that the EUE reliability measures are in closer agreement than the LOLH measures. On a percentage basis, the difference between EUE is about 7%, whereas the LOLH difference is about 12% (see Kahn, 1991). In our judgement, the area under the tail of the outage distribution is likely to be more accurate than its height, as measured by the two models, and this is what we use as the basis for our ELCC calculations.

Table 1. Initial Reliability Indices for Elfin and P+

Model	Expected Unserved Energy (GWh)	Loss-of-load Hours
Elfin	7.4	49.2
P+	6.9	43.9

We chose to maintain as realistic a depiction of the utility as possible, and therefore decided not to adjust loads to an artificial reliability level such as 1 day in 10 years loss-of-load expectation, or equivalent. The ELCC values we calculate are based on those calculated in the base cases illustrated in the table.

For the 100 simulations of the EPA, Figure 7 illustrates the range of ELCC values as a percentage of installed wind capacity. Although these results appear to be consistent with those reported in Milligan (1996a), here we have a larger variance of capacity credit which is because of the larger variation in wind plant output over

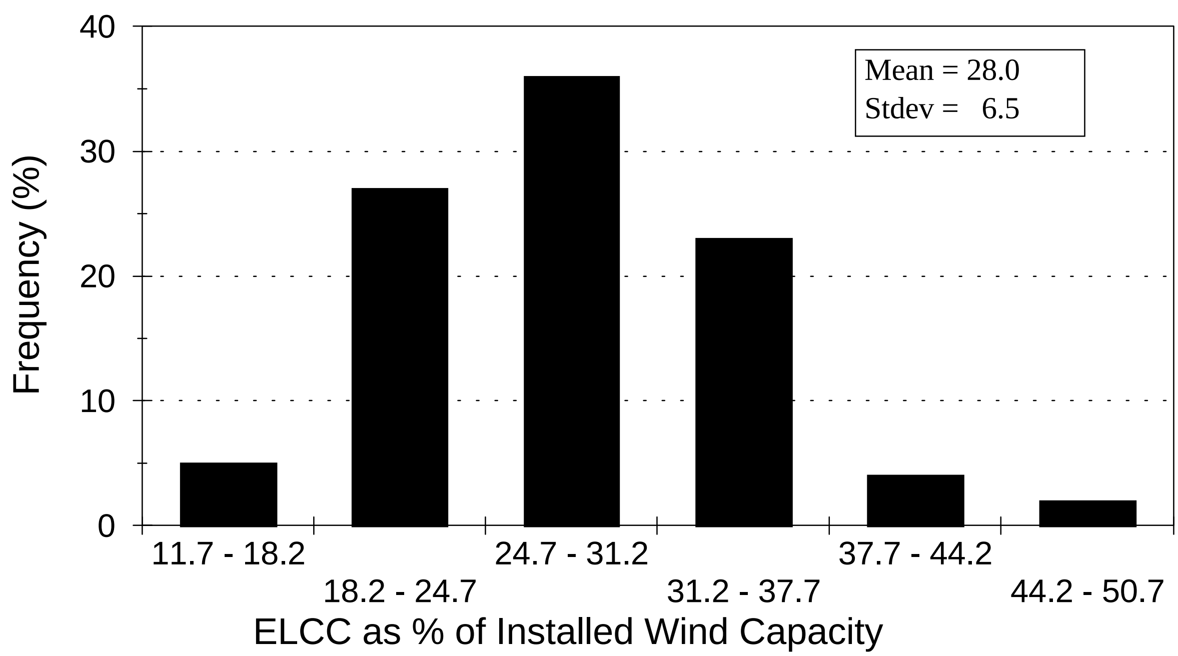

Figure 7. Capacity credit distribution for 100 model runs

the month. The bins for the figure were chosen to be the same size as the sample standard deviation. The graph indicates that all but 2 values are within two standard deviations of the mean.

From the various cases run for the EPA method, we identified those that most closely matched the mean ELCC, and the mean plus or minus the standard deviation of ELCC. The reliability curves for these 3 cases are graphed in Figure 8. The difference between the curves, measured at the 7.4 level of EUE, represents the difference in ELCC between these cases. Figure 9 similarly shows the mean case with the maximum and minimum reliability curves.

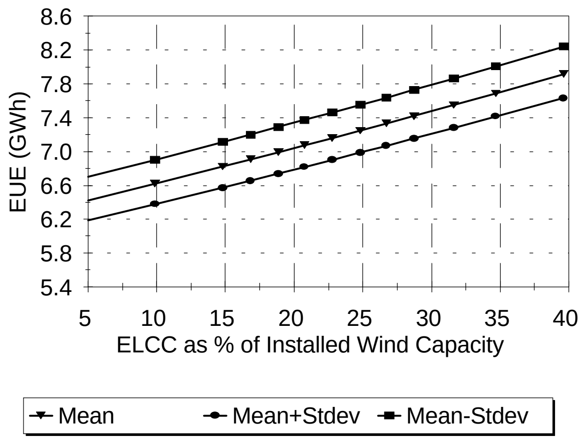

Figure 8. Reliability curves within one standard deviation of the mean ELCC value

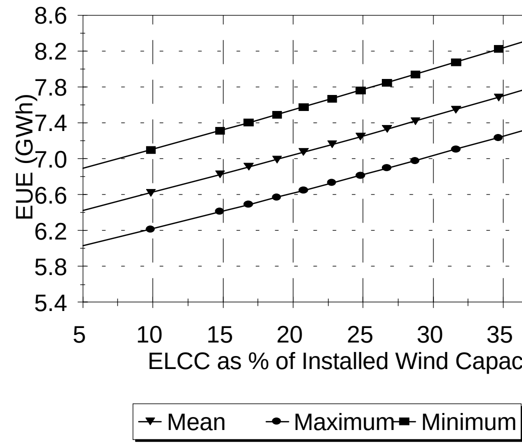

Figure 9. Reliability curves for mean, minimum, and maximum ELCC cases

Figures 10 and 11 show the convergence of the simulations. For each of these graphs, a convergence factor was calculated, which is the ratio of the cumulative mean and standard deviation of ELCC for each iteration. Figure 10 shows the convergence factor, while Figure 11 shows its percentage change. To take full advantage of Monte Carlo simulations such as our EPA, one should be able to specify convergence criteria and run the model until the specified target is reached. Because Elfin is a scenario based model with no intrinsic Monte Carlo capability, we were unable to specify convergence criteria; only the number of runs to perform. See Marnay and Strauss (1990) for further discussion.

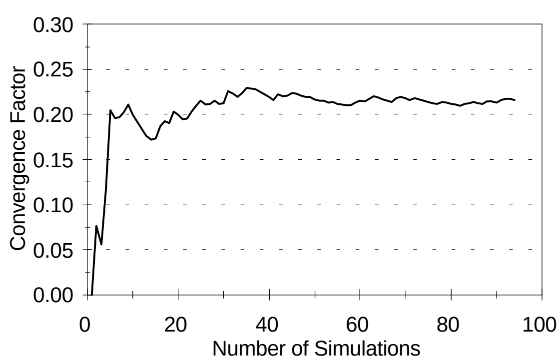

Figure 10. Convergence of the simulations

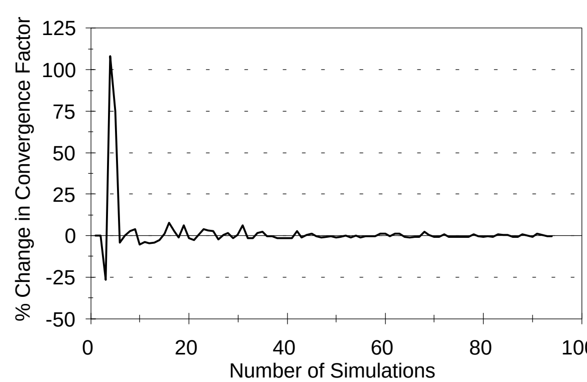

Figure 11. Percent change in convergence factor

The results of the various ELCC calculations are presented in Figure 12. As the figure indicates, there appears to be a closer correlation between the unweighted ELCC from the REPA and EPA than between the weighted REPA and EPA. The chronological model's weighted REPA appears to do a better job than the LDC model's weighted REPA. However, the standard deviation of the EPA and REPA are very close, as shown in Table 2.

Table 2. Comparison of EPA and REPA Means and Standard Deviations

Statistic	Reduced Number of Cases (REPA)	All 100 Cases (EPA)
Mean	29.2	28.0
Std. Dev.	6.4	6.5

The table indicates that the representative cases we selected did retain the variation of the larger sample, as indicated by the standard deviations. The unweighted mean also appears to do a better job than does the weighted mean of estimating that of the larger sample. The algorithm for correctly selecting the representative cases and their respective weights clearly needs additional attention. It is also clear that a selection process applied to a REPA process similar to ours *does* result in a similar estimate of variation of ELCC as does the full EPA case. This implies a considerable savings in both model set-up and execution time, with a small loss of accuracy when compared to the EPA case.

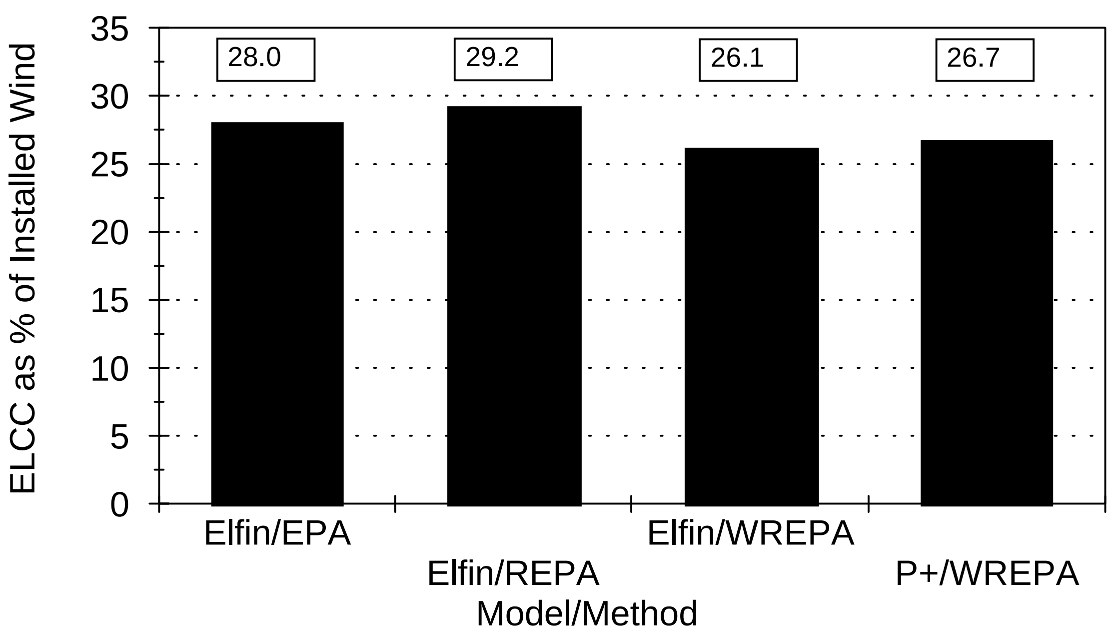

Figure 12. ELCC results

It is also important not to lose sight of our objective: we want to develop a computationally efficient way to provide plausible estimates of wind plant output and *variation* in output. To that end, we believe that the REPA method has accomplished our goal. Further work should be done to explore the impact of a standard method of choosing the data bins so that the subjective element is less of a factor.

Other Selected Results

It is useful to view other results from our model runs. As a chronological model, P+ can produce hourly results for each day. We selected some results of a day in which there was substantial variation to illustrate how our REPA method can be applied to other model outputs of interest. October 14 is the day we chose. The wind power output of the 5 REPA series is reproduced in Figure 13, and Figure 14 shows the average values. It is apparent that there is a great deal of variation on this day that is masked by the average values.

Figure 15 shows the 5 cases (representing the 5 selected bins) that we used for the REPA runs. The figure shows the reduction in base and intermediate generation that is caused by the various wind power series. As can be seen in the figure, there is substantial variation in unit loading. For the case represented by bin 3, reduction in base and intermediate generation is significant, and corresponds closely to the daily peak period. However, the case represented by bin 4 shows a minimal impact on base and intermediate generation. Figure 16 shows the two average cases: one calculated with the weights and the other using a non-weighted average.

We developed similar graphs for peaking generation. Figure 17 shows the change in peaking generation for each of the REPA cases, and Figure 18 shows the average values. Figure 17 also illustrates a wide range of potential relative outcomes. However, the MW level of peaking activity is very small in all cases.

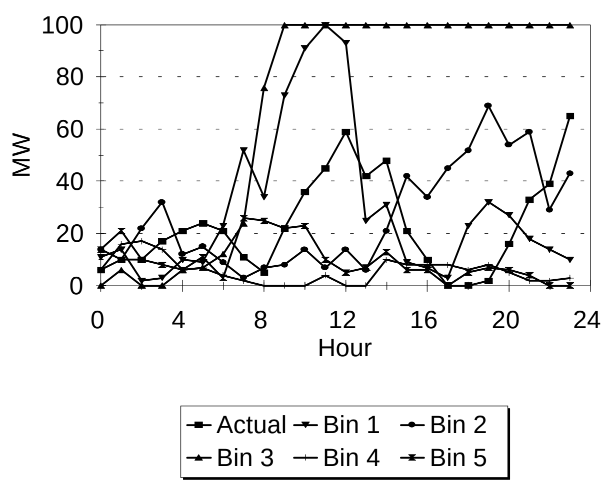

Figure 13. Hourly wind power output, October 14

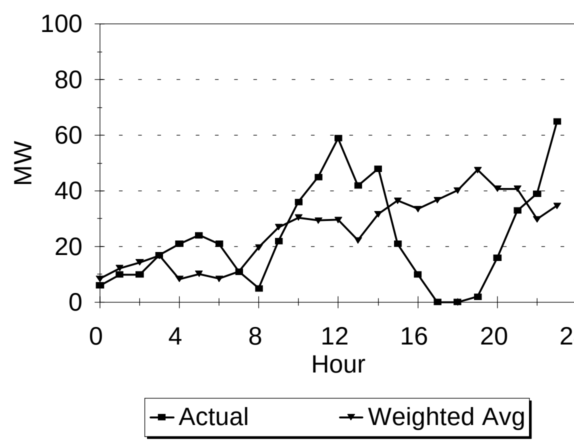

Figure 14. Actual and weighted average of wind power output, October 14

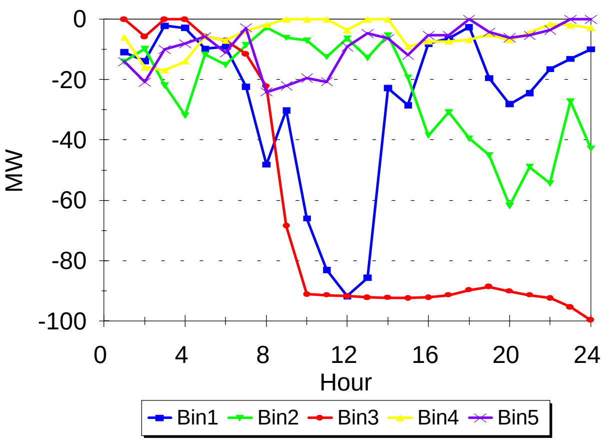

Figure 15. Reduction in base and intermediate generation, REPA cases, October 14

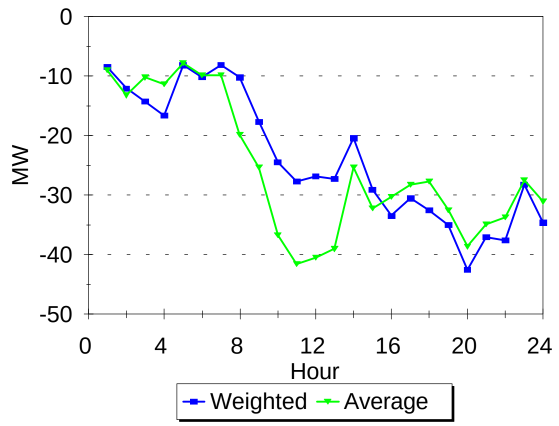

Figure 16. Average reduction in base and intermediate generation, REPA and weighted REPA cases

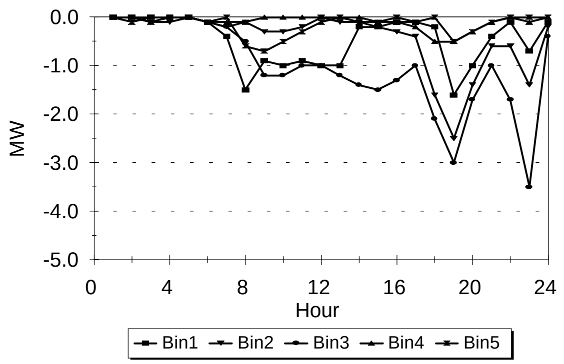

Figure 17. Reduction in peaking generation, REPA cases, October 14

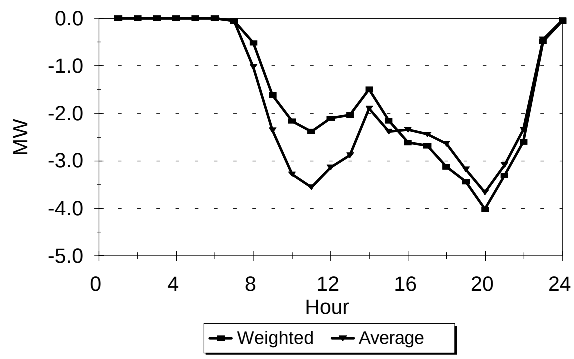

Figure 18. Average reduction in peaking generation, REPA and weighted REPA

We also report some weekly output from the P+ results. In Figure 19, we see the wind-induced change in base and intermediate generation for each of the 5 REPA cases, along with a weighted average for each week. It is apparent from this graph that the weighted average does indeed cover a great deal of variation. Figure 20 is a similar graph of the weekly variation in EUE.

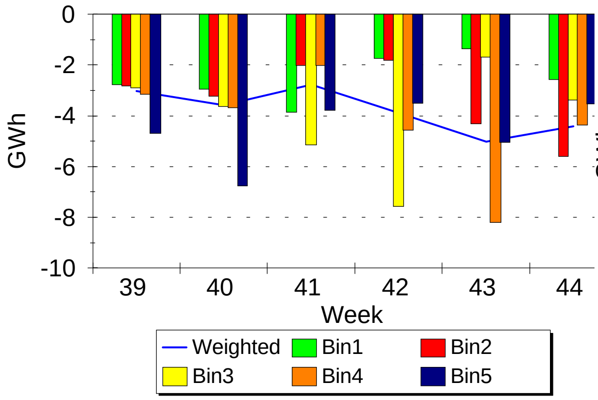

Figure 19. Weekly reduction in base and intermediate generation for 5 REPA cases

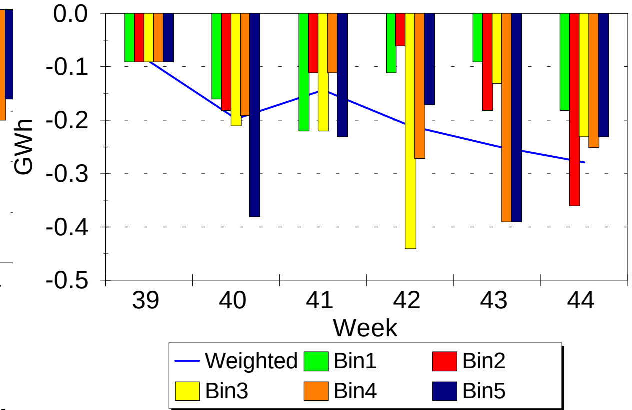

Figure 20. Weekly reduction in EUE for 5 REPA cases

Conclusions

This paper has outlined a computationally efficient way to examine the impact of possible variations in wind plant output. Instead of implementing or modifying a Monte Carlo routine in a production-cost or reliability model, we illustrate a method that can be performed to provide the model with a small number of wind power data sets. The model can then be run for each of the enumerated series, and the results analyzed appropriately for the study at hand. Further refinements can be made in a couple of areas. First, the method of simulating wind data does not have to be Markov, but can consist of any appropriate method. Second, additional experimentation with bin selection could result in a reproducable method that could be converted into a computer algorithm. Although the REPA is not as accurate as the EPA, it does capture the variation in the wind resource.

REFERENCES

Bernow, S., B. Biewald, J. Hall, D. Singh (1994), "Modelling Renewable Electric Resources: A Case Study of Wind." Tellus No. 91-187. Tellus Institute: Boston, Massachussetts.

Billinton, R. and R. Allan (1984). *Reliability Evaluation of Power Systems*. New York: Plenum Press.

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.

Deshmukh, R. G. and R. Ramakumar (1982). "Reliability Analysis of Combined Wind-Electric and Conventional Generation Systems." *Solar Energy.* Vol 28, No. 4.

Garver, L. L. (1966). "Effective Load-Carrying Capability of Generating Units." *IEEE Transactions on Power Apparatus and Systems.* Vol PAS-85.

Hillier, F. and G. Lieberman (1974), *Operations Research.* San Francisco: Holden-Day.

Jarass, L., L. Hoffmann, A. Jarass, G. Obermair (1981), *Wind Energy: An Assessment of the Technical and Economic Potential.* Springer-Verlag. New York.

Kahn, E. (1991). *Electric Utility Planning & Regulation.* American Council for an Energy-Efficient Economy. Washington, D.C.

Manwell, J. F., G. Deng, and J. G. McGowan (1994). "A Markov Process Based Performance Model for Wind/Diesel Battery Storage Systems." *European Wind Energy Association Conference and Exhibition; October 10-14, 1994;* Thessaloniki - Macedonia, Greece. Roskilde, Denmark: Department of Meteorology and Wind Energy, Risö National Laboratory.

Marnay, C. and T. Strauss (1990). *Variance Reduction in Monte Carlo Chronological Production-cost Modeling.* Sacramento, California: California Public Utilities Commission.

Milligan, M. (1996a). "Variance Estimates of Wind Plant Capacity Credit." *Windpower '96 Proceedings; June 23-27, 1996;* Denver, Colorado. NREL/TP-440-21311. Golden, Colorado: National Renewable Energy Laboratory.

Milligan, M. (1996b). *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. 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.

Utility Wind-Modeling Planning Meeting. (1996). National Renewable Energy Laboratory, Golden, CO, 22-23 February.