

Modeling Utility-Scale Wind Power Plants Part 2: Capacity Credit

Michael R. Milligan



NREL

National Renewable Energy Laboratory

1617 Cole Boulevard
Golden, Colorado 80401-3393

NREL is a U.S. Department of Energy Laboratory
Operated by Midwest Research Institute • Battelle • Bechtel

Contract No. DE-AC36-99-GO10337

Modeling Utility-Scale Wind Power Plants Part 2: Capacity Credit

Michael R. Milligan

Prepared under Task No. WER1.3010



NREL

National Renewable Energy Laboratory

1617 Cole Boulevard
Golden, Colorado 80401-3393

NREL is a U.S. Department of Energy Laboratory
Operated by Midwest Research Institute • Battelle • Bechtel

Contract No. DE-AC36-99-GO10337

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 electronically at <http://www.osti.gov/bridge>

Available for a processing fee to U.S. Department of Energy
and its contractors, in paper, from:

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
phone: 865.576.8401
fax: 865.576.5728
email: reports@adonis.osti.gov

Available for sale to the public, in paper, from:

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
phone: 800.553.6847
fax: 703.605.6900
email: orders@ntis.fedworld.gov
online ordering: <http://www.ntis.gov/ordering.htm>



Abstract

As the worldwide use of wind turbine generators in utility-scale applications continues to increase, it will become increasingly important to assess the economic and reliability impact of these intermittent resources. Although the utility industry appears to be moving towards a restructured environment, basic economic and reliability issues will continue to be relevant to companies involved with electricity generation. This paper is the second in a two-part series that addresses modeling approaches and results that were obtained in several case studies and research projects at the National Renewable Energy Laboratory (NREL). This second paper focuses on wind plant capacity credit as measured with power system reliability indices. Reliability-based methods of measuring capacity credit are compared with wind plant capacity factor. The relationship between capacity-credit and accurate wind forecasting is also explored.

Table of Contents

Introduction	1
Capacity Credit.....	2
Capacity Credit in an Operational Setting	2
Wind Plant Capacity Credit and Wind-Speed Forecasting	4
The Value of Wind Forecasting.....	4
Modeling the Wind Forecast Scenarios	6
Forecast Results	8
Forecast Benefits Summary	10
Wind Plant Capacity Modeling.....	10
Methods for Calculating Capacity Credit	12
Utility Production and Reliability Modeling.....	12
Effective Load-Carrying Capability.....	12
Application to Intermittent Resources	14
Simpler Approaches	15
Application of Alternative Methods	15
Results	16
Comparison of Methods	18
Summary.....	19
Sequential Monte Carlo.....	21
A Markov Model for Wind-Speed Variations	22
Capacity Credit Modeling with Elfin: Comparison of the Markov Model to Actual Wind-Speed Data..	24
Capacity Credit Assessment Using Sequential Monte Carlo: P+ and Elfin	26
Capacity Credit Results	29
Summary.....	32
Modeling Production-Cost Variations with Sequential Monte Carlo.....	33
Summary.....	38
Modeling Geographically Disperse Wind Power Plants.....	39
The Potential Benefit of Wind Power Plant Geographic Diversity.....	39
State of Minnesota Data Collection Project.....	39
Wind Site Selection	40
Modeling	43
Marginal Reliability Method.....	44
Optimization with Elfin	46
Inter-Annual Variations and Uncertainty.....	47
Comparison of Results.....	48

Summary.....	49
Conclusions.....	50
Modeling Results.....	50
Modeling Issues.....	51
Assessing the Modeling Frameworks	52
Future Research	53
References.....	55

List of Figures

Figure 1. Hypothetical wind plant operational capacity credit using 1-year actual load and wind data.....	3
Figure 2. Hypothetical wind plant operational capacity credit using 10-year actual load.....	3
Figure 3. Effect of forecast error on benefit, Utility A, 50% output.....	9
Figure 4. Effect of forecast error on benefit, Utility A, full output.....	9
Figure 5. Reliability curves for calculating the ELCC of a wind plant.....	14
Figure 6. Capacity factor for top hours, various methods, 1980.....	16
Figure 7. Capacity factor for top hours, various methods, 1982.....	17
Figure 8. Capacity factor for top hours, various methods, 1984.....	17
Figure 9. Capacity factor for top hours, various methods, 1990.....	18
Figure 10. Overall comparison of methods	19
Figure 11. Peak month state transition matrix	23
Figure 12. Capacity credit and energy for 13 years	24
Figure 13. Comparison between simulated and actual ELCC	25
Figure 14. Comparison between simulated and actual energy.....	26
Figure 15. Wind-speed state transition matrix for October	27
Figure 16. Wind energy distribution for multiple data sets, October.....	28
Figure 17. Wind energy distribution, binned by standard deviation.....	29
Figure 18. Capacity credit distribution of EPA runs	30
Figure 19. Capacity credit results from various methods	31
Figure 20. Comparison of the REPA cases by model.....	32
Figure 21. Correlation between wind energy and ELCC.....	32
Figure 22. Generation reduction, week of October 9	33
Figure 23. Generation reduction, October 9.....	34
Figure 24. Change in base and intermediate generation from no-wind case, REPA cases, October 14	35
Figure 25. Change in base and intermediate generation, mean REPA and weighted REPA cases, October 14.....	35
Figure 26. Change in peaking generation, REPA cases, October 14.....	36
Figure 27. Change in peaking generation, mean REPA and weighted REPA cases, October 14.....	36
Figure 28. Weekly change in base and intermediate generation for 5 REPA cases.....	37
Figure 29. Weekly change in EENS for 5 REPA cases.....	37
Figure 30. Minnesota wind sites.....	41
Figure 31. Wind power output over a 2-day period from multiple sites.....	42
Figure 32. Wind power output over a 1-day period from multiple sites.....	42
Figure 33. Minnesota annual load-duration curve	43
Figure 34. Marginal reliability curve for Becker	45

Figure 35. Marginal reliability curve for Brewster	45
Figure 36. Optimization algorithm	46
Figure 37. Comparison of methods	49

Introduction

Capacity credit estimates of wind power plants help generating companies, utility planners, and other decision-makers evaluate this intermittent resource in the context of other types of power plants. Capacity credit is the level of conventional generation that can be replaced with wind generation. To perform such an analysis, it is important to define the way in which one type of resource can be substituted for another. Most analysts prefer basing such a trade-off on a reliability measure. A common measure of system reliability is loss-of-load expectation, or LOLE. The LOLE is an indication of the statistically expected number of times within a given time period that the system could not provide for customer load. If a given level of wind-generating capacity can be substituted for conventional capacity, holding the reliability level constant, then we can obtain a measure of wind plant capacity credit.

In Part 1 [1], we examined the overall economics of wind power plants by using production-cost models to estimate the benefits. The focus of that work is the wind energy displacement value, but other factors were considered, albeit briefly. In Part 2, the focus is on reliability-based measures of wind power plants. However, we will also look at operational capacity credits that are based on two power pool rules, and consider how to maximize the reliability of wind power plants at geographically disperse locations. The results in this paper are based on hourly data. NREL is undertaking a joint research project to analyze the impact that wind power plants have on ancillary services, based on wind-power data at time resolutions down to 1 second. The results from this study are anticipated sometime in 2002. As restructuring continues to lurch forward, we continue to use the term “utility” in a somewhat generic sense, applying this term to both the electric utilities of regulated markets and generating companies in restructured markets. Although there are clear differences in the structure of these companies, it is beyond the scope of this paper to analyze these differences.

Capacity Credit

Utilities are often reluctant to assign any capacity credit to wind power plants, largely because of the intermittent nature of the resource and the perceived difficulties in accurately forecasting wind power output. In fact, the capacity value of a wind plant can have several meanings, depending on the context in which the discussion takes place. Local reliability councils frequently have unit accreditation standards, which normally cover conventional generators. These capacity accreditations are operationally based, are designed to ensure the power system's reliable operation, and are normally based on demonstrations of the plant's continuous operation for several hours. (Although transmission and distribution systems also play an important role in overall reliability assessment, this paper will focus only on generator reliability.) In a recent NREL study, only two North American Electrical Reliability Council regions were found to explicitly cover accreditation of wind power plants [2].

This operational accreditation of conventional power plants is not normally based on a reliability measure, such as effective load-carrying capability (ELCC). ELCC is a way to measure a power plant's capacity contributions based on its influence on overall system reliability. Using a measure such as ELCC, all power plants with a non-zero forced outage rate have an ELCC that is less than rated capacity (barring unusual plants with artificially low-rated capacity with respect to actual achieved capacity). The ELCC measure is often used as a way to compare alternative power plants, and can be easily applied to wind power plants as well. A power plant's ELCC is typically calculated with an electric system reliability model or by a production-cost model.

This section discusses operational capacity credit, the impact of accurate wind forecasting on unit commitment, and reliability-based capacity credit using Elfin. Additional analysis of wind plant capacity credit can be found in the Sequential Monte Carlo section, below.

Capacity Credit in an Operational Setting

As described by Wan, most reliability councils do not have specific procedures for rating wind power plants. If a wind power plant cannot claim any operational capacity credit, then other generating resources must be committed in an amount equal to the operating level for the wind plant for a specific time period. A wind power plant will have more value if it can replace conventional committed capacity, at least for some portion of the year. A generating company that owns significant wind capacity and a slow-start generator, such as coal, will have an incentive to minimize start-ups of the coal plant, because each start-up has a fixed cost. If wind capacity can be counted on as "firm" for several hours, a potentially expensive thermal plant start-up might be avoided. This is balanced against the need for a reliable generating system. If we were to avoid starting a conventional generator during a period that sufficient wind capacity is expected, and if the wind capacity did not materialize, system reliability would be compromised.

Wan provides some numerical examples of operational capacity credit, as calculated by the Mid-Continent Area Power Pool (MAPP) and New England Power Pool (NEPOOL) procedures. Each method is applied to one year of wind data (the minimum allowable for each procedure) and ten years (recommended) of wind data. The results of Wan's analysis are shown in Figures 1 and 2. We can conclude from these graphs that the two rating methods are not all that different, and that wind plants do indeed have capacity credit in

an operational context. We can also see the effect of using more data in the calculation, which also tends to minimize differences between the methods. However, it is not clear that these methods provide a scheduler with the information needed to optimize the commitment of conventional plants. For that, we turn to predictions of the future.

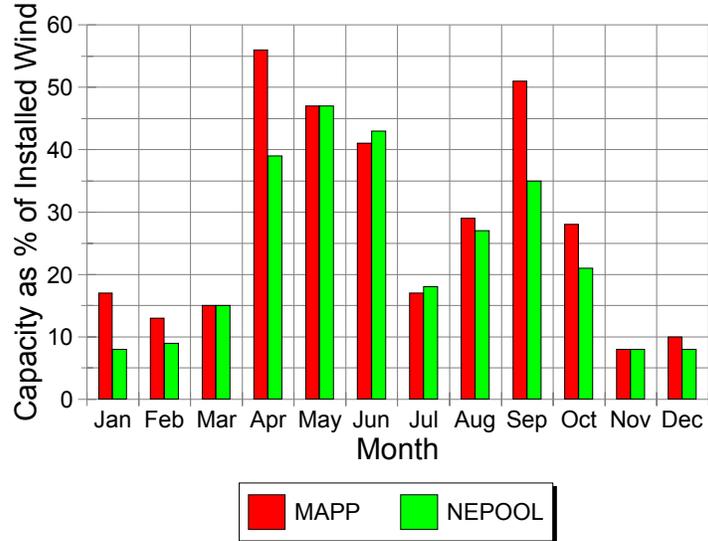


Figure 1. Hypothetical wind plant operational capacity credit using 1-year actual load and wind data

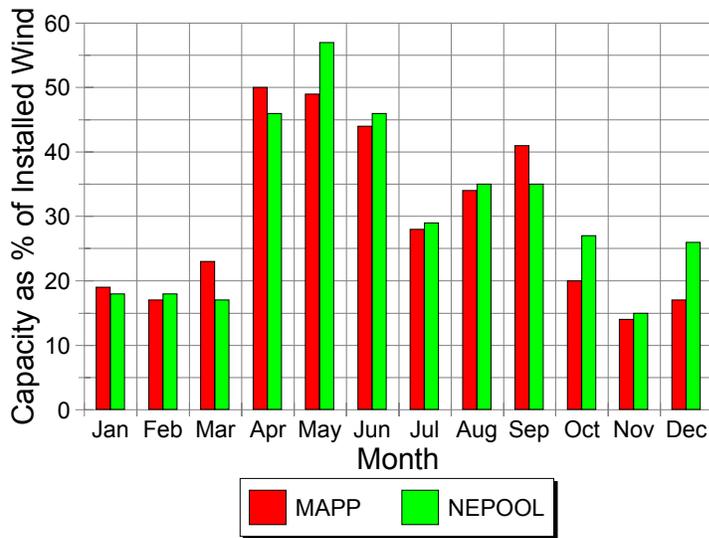


Figure 2. Hypothetical wind plant operational capacity credit using 10-year actual load and wind data

Wind Plant Capacity Credit and Wind-Speed Forecasting

In an operational setting, the hourly capacity of a wind power plant can be bid into a pool if the wind plant's output can be accurately predicted 24 to 48 hours in advance. In states where restructuring is still under investigation, the unit commitment decision is also made a day or so in advance. Knowing the level of wind power output in advance can help in both situations. Milligan, Miller, and Chapman [3] showed significant economic benefits of accurate wind forecasts because the unit commitment decision can effectively use an accurate wind power forecast throughout the day. NREL is currently working with the Electric Power Research Institute on a wind-energy-forecasting development and testing program, and is also conducting independent research on wind forecasting techniques. In the United States, commercial wind forecasting services are available, and European forecasting models are now being successfully applied to several wind power plants [4]. Although accurate wind forecasts help with the unit commitment problem, they may also significantly reduce energy imbalance charges.

The simplest benefit of an accurate wind forecast is that wind-generated energy can be planned for and used by the utility so the utility avoids the need to consume fuel to produce electricity. Wind can be viewed as a free fuel. If all of the wind turbine's output replaces the output of generators that consume the most expensive fuel, then savings are maximized. This simplified point of view is complicated by constraints imposed by integrating the wind resource with the rest of the electricity supply system.

Wind power has value even if forecasts are not accurate. Clearly, the amount of wind energy available at a given time is independent of the forecast. Unexpected wind power can still be used to displace energy that is provided by a load-following unit. However, as wind plants continue to become more economically competitive with conventional energy sources, it is important to correctly assess the capacity value of wind. A wind resource that cannot be counted on for capacity has a minimal contribution to system reliability. An accurate forecast allows the utility to count wind capacity and reduce costs without violating reliability constraints.

Both commitment and spinning reserve must be considered by utilities when scheduling resources to meet load, and both affect the value of a wind forecast. Understanding these two processes is helped by a quick outline of some issues affecting utility resource planning decisions.

The Value of Wind Forecasting

This discussion follows Milligan, Miller, and Chapman, whose goal was to quantify the potential value of an accurate wind forecast. Although the original research was based on a regulated environment, the results have a bearing on restructured markets, or on any electricity market in which buyers and sellers contract for electricity in day-ahead or several-day-ahead markets. To assess this value, the Elfin model was applied to utilities A and B (the utility names are not revealed because of prior agreement), and commitment errors were simulated in such a way as to approximate wind forecast errors. The economic value of an accurate wind forecast is based on the utility or generating company's ability to perform optimal unit commitment based on the forecast. This discussion begins by showing the relationship between the wind forecast and unit commitment decisions, along with implications for various types of power plants. This is followed by some modeling results.

An electric utility or ISO meets instantaneous demand by coordinating output from a variety of generation sources. Multiple independent generator output must be provided in such a way as to provide electricity to the distribution grid with the correct frequency and voltage. If a generator fails suddenly, and there is no compensating action to make up for the lost capacity, the whole grid can fall out of balance, potentially requiring a system-wide shutdown and restart. To avoid these difficulties, a utility will keep some capacity spinning in reserve, (i.e., a running generator capable of producing electricity but not providing current to the grid). If an on-line generator fails, the spinning capacity contributes to the grid and restores balance.

Although wind turbines are intermittent generators, it is not generally necessary to provide standby capacity that specifically provides backup to the wind power plant. The principle of diversity and the low probability of simultaneous outages from multiple power plants have led to calculating spinning reserve requirements as a function of load or largest contingency. (NREL is currently developing a method that allocates the system-wide spinning reserve requirement to each power plant in the system, based on its capacity and forced outage rate [5].) However, deviations from the forecasted level of wind power output could have an impact on reserves, which implies that the forecast could be beneficial if it prevents a drop in reserves during unanticipated drops in wind power output.

Commitment issues involve a longer time scale than spinning reserve issues. The amount of lead time required for a generator to go from a non-operating state to an electricity-producing state varies among technologies. An example of a “quick-start” unit is a combustion turbine, whose lead time can be measured in minutes. An example of a “slow-start” unit is a nuclear plant, which requires hours or days of lead time to bring to an operating state. Once a slow-start unit is operating, it must continue to run or go through a lengthy shutdown and restart. We refer to a slow-start unit’s generation threshold, the level below which the unit must be shutdown and restarted, as its minimum output level. Slow-start units are ramped down during low-load periods, possibly as far as their minimum operating level, and then ramped up to assist in meeting peaks. Intermediate and peaking units are also ramped up or started, or both, to meet the peak. This method of operation is called load following.

An electric utility meets loads through a combination of slow- and quick-start capacity. Typically, the utility must plan at least one day ahead to have sufficient slow-start capacity online to be able to meet the next daily peak. The process of deciding which slow-start units to operate is known as the unit commitment problem.

The utility must commit sufficient slow-start capacity to be able to meet as much demand as possible with low-cost generators. However, committing too much slow-start capacity has a cost. When the system reaches minimum load conditions, the capacity across the minimum operating levels of committed slow-start plants may exceed demand, causing the utility to dump or sell excess energy, perhaps at a price below marginal cost. Utilities are constrained by the lead times required to cool down and restart slow-start units; they cannot simply turn off slow-start units that are surplus at night and restart them for the next day's peak. Typically, a utility avoids the problem of having too much slow-start capacity during non-peak hours by relying on quick-start units to serve peak loads for a few hours each weekday.

In contrast to over-commitment problems, if a utility under-commits (i.e., starts up too few slow-start units), it will rely excessively on higher cost quick-start energy or purchases to meet peaks. In the extreme case, the utility may even be short of capacity at the time of system peak.

Optimized commitment is complicated by intermittent output from a wind resource. If wind generation represents a significant portion of the total capacity, the utility must balance the risk of the wind resource being a non-contributor against the risk of committing excess slow-start capacity. An accurate wind forecast would allow the utility to optimize the amount of slow-start capacity it commits, thus minimizing the risk of producing either too much or too little low-cost energy.

An accurate forecast could enable a utility to plan its generation schedule by using the wind power forecast as the “schedule” of wind output. As long as the forecast is accurate, the utility is in an optimal position, able to integrate wind output without backup or support for the wind resource. An inaccurate forecast, however, causes problems. If the forecast is too optimistic (i.e., forecasted wind exceeds actual wind), then the utility finds its reserve margin too low. A capacity deficit could violate minimum reliability constraints and may cause system costs to be higher than necessary.

Conversely, the reliability target can be exceeded when the forecast is too pessimistic (the forecast wind is less than the actual wind). In this case, the utility has completed its commitment list, assuming the forecast wind power level. The actual wind power output is greater than forecast. From a reliability perspective, units are committed that need not be, resulting in excess spinning capacity. While the excess reliability is not a problem itself, it does result in a non-optimal schedule of resources and needlessly increases energy costs.

Modeling the Wind Forecast Scenarios

We modeled differing degrees of accuracy in a wind forecast and measured the total cost of producing electricity in each scenario. The differences in total cost for each scenario were measured and used to project the benefits of accuracy in the forecast. All production simulations were carried out for a single year, and forecast benefits are measured as annual values. There are long-term implications that wind-forecasting accuracy might have on optimal generation expansion, but they are not explored here.

We developed two methods to analyze the benefit of a wind forecast. The first method, which uses hourly wind power data, is described here. Details on a second method can be found in the original paper. The scenario begins with the meteorological forecast, which results in an hourly wind-speed forecast for the time period in question (on the order of 24-48 hours). These data are used to calculate the hourly wind power output, which is then given to the utility dispatcher. The power system schedule is then based on the wind forecast, and decisions are made on which of the conventional units should be committed. Actual wind power may vary from its predicted value. The extent of the wind forecast error determines the financial and reliability cost to the utility. If the forecast error is zero, then the utility has planned in an optimal way. Penalties occur when errors in forecasting occur, and these penalties increase as a function of the forecasting error. The modeling attempts to capture these scenarios in as much detail as possible. We provided Elfin with the hourly wind power forecast, so the model can complete optimal resource scheduling. The wind forecast is treated as any other scheduled resource and is considered firm. We then provide Elfin with hourly deviations from the forecast. To capture the best time resolution possible, the wind resource was modeled as a load modifier. While it would be useful to use the probabilistic modeling technique discussed above, it was not possible to model forecast errors with that method.

To account for the effects of wind power deviating from the forecast, we assume the utility relies on the forecast in planning its commitment schedule. Wind power above or below the forecast is not counted towards the commitment target. In other words, the forecast wind power counts towards commitment, whereas the deviation from the forecast does not. Among all scenarios with a constant amount of wind

energy, differences in total production cost are a function of the portion of wind output that was counted on during the commitment process.

This method is extremely flexible and is a reasonable way to approximate the forecasting error. The method allows us to separately examine the question of the effects of over-forecasting versus under-forecasting. As examples, consider these cases. First, suppose that the forecast is completely accurate. That implies that the non-firm component of wind is zero, and the conventional generators can be scheduled in an optimal way (ignoring other sources of uncertainty). If the forecast is too high, we reduce actual wind power from the forecast level, but do not allow commitment to change. Similarly, if the forecast is too low, we increase available wind power, but do not count this increase towards commitment.

This method could also be expanded so that other forecast error patterns can be introduced. Here, we assume the forecast error is consistently higher or lower than actual wind output. In reality, forecast errors are likely to be too high in some hours and too low in others. It is also important to note that, unless forecast error is explicitly taken into account, it is *implicitly* assumed to be zero.

For this study, we did not allow Elfin to price any resources at system lambda, because altering the forecast error tends to change system marginal costs and artificially alters total cost. They also held the load profile constant, so that differences in the correlation between wind power and load would not influence the results.

We used a single year of wind data from the High Plains wind site, calculating equivalent power output, and running the Elfin model. We ran 112 different cases per utility, each of which represents a different combination of actual and forecast wind power. The actual wind power was simulated at 10% increments of the historical wind levels. If, for example, the historical wind data resulted in wind output of 100 MW for the hour, the model was run at 10 MW, increasing by 10 MW, up to 100 MW. For each of these wind power levels, there are several ways of combining 10% increments of forecast wind and forecast error values. For example, 50 MW of actual output can be realized by several combinations, such as a 60-MW forecast and -10-MW forecast error; a 70-MW forecast and -20-MW forecast error, or a 30-MW forecast and 20-MW forecast error. For the site maximum wind speed, which translates to rated wind plant output, the forecast error can only be positive. Likewise, for 0 MW actual wind power, the forecast error cannot exceed 0. However, zero wind output can be achieved with offsetting values of forecast and forecast error: 0 forecast and 0 error; 10 forecast and -10 error, for example.

The method allows us to choose any combination of wind power forecast and forecast accuracy. For example, during a given hour, the forecast may be for 50 MW of wind, and the actual wind output might be 80 MW—the forecast error is therefore 30 MW. In an extreme case of a missed forecast, a 100-MW forecast can be simulated with 0 MW wind (forecast error is 100 MW). In this case, the utility has planned on 100 MW of wind, but must dip into its reserves to meet the load. Conversely, the utility can plan on 0 MW of wind and receive 100 MW. However, spinning reserve and commitment levels are set in advance, so the benefit of 100 MW of unplanned wind power is diminished. It is important to note that wind energy is held constant to assess the value of the forecast. Therefore, the primary impact of the forecast error enters via the unit commitment logic. It is also important to note that spinning reserve targets were also held constant in the various scenarios. However, the units that are called on to meet the fixed spinning reserve target may, in fact, differ from the perfect forecast case because of differences in unit commitment that arise from differing forecast accuracies.

To capture commitment effects, the wind resource had to be sizeable enough to cause changes in total production cost because of changes in the wind plant's firmness, or reliability. Therefore, the capacity portion was modeled to yield a constant 4% of effective system peak-coincident firm capacity across all cases (i.e., 4% of the peak load would be expected to be served by the wind plant). By modeling varying degrees of forecast accuracy and recording the total cost at the optimal commitment level—holding the wind energy contribution and all other factors constant—we calculated the change in system variable costs that were caused only by the change in the forecast accuracy.

Forecast Results

Some of these results are presented in Figures 3 and 4. All the cases show that an accurate wind forecast does give the most benefit from the wind resource and that inaccurate forecasts reduce benefits. The direction of the forecast error does influence the benefit reduction for a given utility, but is not constant between the utilities in this study. The relative costs of over- or under-forecasting will depend on the tradeoffs faced by each utility (e.g., what is the penalty imposed by the power pool or regulators for under-committing, as balanced against the cost of over-committing).

This work shows that an accurate wind forecast does have value. Several conclusions can be drawn, although the specific results depend heavily on the utility's generator mix, load swings, correlation of wind to load, and other factors. First, accurate forecasts always have economic value, even at lower levels of wind output. Second, even if there is no wind power available, it is in the utility's best interest to forecast this event accurately. Third, the results indicate that a forecast that is too high is more costly than one that is too low for Utility A, but the opposite is true for Utility B. These results are highly sensitive to the utility's contractual and power pool arrangements and should be studied more closely in that context.

Figure 3 assumes that actual wind power output is at 50% of rated capacity for the wind plant. The benefit, as measured on the vertical axis, shows that with a perfectly accurate forecast (forecast error of zero), there is about a \$54 million benefit provided by the wind plant. This benefit is measured as fuel saving from conventional units. The diagram shows how this benefit is reduced as the absolute value of the forecast error increases. In the case of a 50% forecast error, the benefit of the wind plant is reduced to about \$50 million. Therefore, the forecast provides approximately \$4 million in benefits. Conversely, a negative 50% forecast error results in a benefit reduction of about \$2 million. The graph in this figure is a typical shape for the various percentages of wind capacity that were modeled.

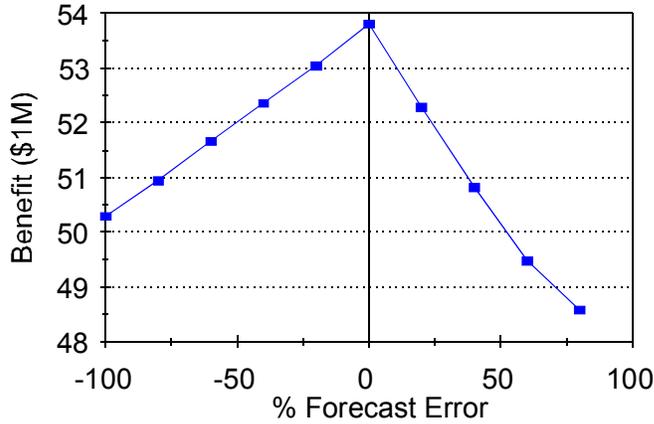


Figure 3. Effect of forecast error on benefit, Utility A, 50% output

Figure 4 assumes full wind plant output. With full output, it is not possible for the forecast to exceed the actual wind. Therefore, the graph only shows positive forecast errors (i.e., in all cases, actual output exceeds forecast output). This graph shows that a forecast error of 90% costs the utility about \$12 million. The consequences of less severe errors can be read from the graph—for example, a 40% forecast error results in approximately \$7.5 million in lost benefits.

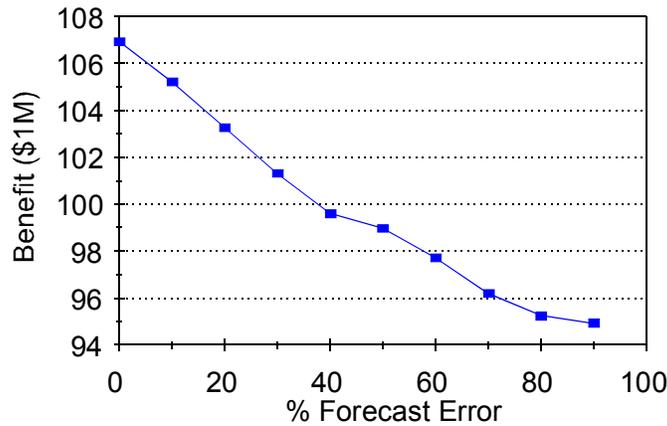


Figure 4. Effect of forecast error on benefit, Utility A, full output

The test wind plant has a capacity of 1,250 MW. After accounting for various losses, the net output capability is 1,006 MW. For such a wind plant, the results indicate that an accurate forecast can be worth up to about \$12 million per year for utility A, and up to \$40 million per year for utility B. The difference in results for the two utilities reflects the differences in costs of the marginal units displaced by wind generation. A utility that has a relatively large slow-start unit on the margin during low-load periods is likely to see a larger cost impact from an accurate forecast. Of course, it is extremely unlikely that forecasting with 100% accuracy can be achieved on average, but this does suggest that it would be worthwhile for a utility to pursue some reasonable level of forecasting accuracy.

Forecast Benefits Summary

Forecast benefits will vary according to the utility's load pattern, generation mix, fuel costs, correlation of load with wind power, and other factors. However, in all the cases analyzed, an accurate wind forecast has economic benefits. The declining marginal benefit curve from the forced-outage method suggests that the goal of a wind forecasting project should not be 100% accuracy, but should balance the degree of accuracy against the degree of benefit derived by that level of accuracy. However, further work in this area is clearly needed.

This method allows us to distinguish between the impacts of under-forecasting and over-forecasting wind power. As outlined in the discussion above, the reliability and economic impacts of under- or over-forecasting may not be symmetric because of differences in the cost structure and consequences of over-committing and under-committing. Furthermore, power companies typically have marginal cost curves that increase with output, which also implies that the cost of under-forecasting would likely exceed the cost of over-forecasting. The benefits of accurate wind forecasting that were calculated are significant, but should be interpreted as a first approximation of an iterative process. Further improvements in techniques and using an actual wind-forecasting model could improve the accuracy of these results. It is also important to note that, given an inaccurate forecast, the only effect of the forecast error is via the utility's commitment decisions. Therefore, we assume that all wind energy is used by the utility. This means that the economic benefits measured here may be understated. However, given these caveats, a couple of points can be made. First, a utility that treats wind generators purely as economical energy resources with no firm capacity would be the best candidate to invest in some degree of accuracy in a wind forecast. Second, concentrating wind generators in regions where wind forecasts were maximized for accuracy could bring economies of scale, resulting in more cost-effective forecasting.

Third, utilities are already managing uncertainty and intermittence in forecasting load, even without using wind or other intermittent power generators. Weather forecasting plays a role in load forecasting, especially in areas that experience extreme weather conditions. Many utilities have already invested in weather forecasting and may be able to leverage existing weather forecasting capabilities in developing wind-forecasting expertise. Other intermittent, weather-dependent renewable resources, such as solar-powered generators, could also benefit from accurate forecasting techniques.

To the extent that an improved wind forecast reduces the use of fossil fuels, pollution would also be reduced. If this cost were internalized by the utility, the economic benefit of the accurate wind forecast would be increased accordingly.

Wind Plant Capacity Modeling

As electricity markets evolve towards a more competitive structure, the traditional vertical integration of electric utilities is evolving into separate business units for electricity generation, transmission, and distribution. Although the precise form of these new industries is not known for certain, it is clear that the generating industry will continue to evaluate whether or not to build new generators. The companies that will perform this type of analysis will probably be very diverse, and may include very large generating companies that have resulted from the many recent mergers and acquisitions, as well as smaller companies that are currently considered independent power producers. In my judgement, some entity in the brave new world will continue to be interested in ELCC or other reliability measures to ensure adequacy of the generating system [6].

Utilities have historically entered into agreements with each other, buying and selling electricity on the wholesale market. The structure of these contracts depends on the needs of the parties involved. In some cases, payment is based only on the quantity of electric energy purchased during the contract period. These are typically called “energy-only” contracts. In other cases, however, the payment is based, at least in part, on the capacity that is provided. This type of agreement is generally called a “capacity-only” contract. To thoroughly evaluate the return on investment, investors must perform calculations that allow them to estimate long-term capacity sales, whether to retail customers or wholesale customers. In the restructured markets of the future, it is not clear whether market players will have the technical capability to perform the capacity and reliability analysis that has historically been done by the traditional utility. Various other methods have been proposed (see [7] and [8] for discussions) that involve using the capacity factor to estimate the long-term capacity credit for a wind generating plant.

Generally, energy-only contracts are based on incremental costs such as fuel usage and operations and maintenance (O&M) costs associated with the increased use of the power plant. However, renewable energy resources, such as wind and solar use “free” fuel, and have low O&M costs. The owner of a wind plant who could sell only energy competes with the fuel cost of other plants. These payments must be sufficient over the life of the wind plant to recover the capital. Currently, incremental energy and O&M costs are very low, so, as discussed in a previous section, if a solar or wind resource is predictable or reliable enough, the potential for increased payment based on capacity is important.

Another difficulty arises because of the intermittent nature of renewable resources such as wind and solar. During windy periods, fuel is available to drive the wind plant. During lulls, fuel (i.e., the wind) is not available, and wind plant output will fall, possibly to zero. If a wind plant operator should enter into a capacity sale agreement during a period in which expected wind output did not occur, there would be a financial penalty to the seller. The purchaser may not have sufficient resources, and would perhaps be unable to meet customer loads. Additional purchases or generation would be required to make up the deficiency. Conversely, if the wind operator were to underestimate the capacity available during a particular time period, the excess capacity would go unused and unsold, reducing the payments to the wind plant. In this case, the purchaser would have greater capacity than needed. In cases of high excess capacity, conventional generators might need to back off, providing a bonus in fuel saving.

The need for accurate capacity assessment is clear. However, the traditional approach involves using complex electricity-production simulation models, which may incur a significant cost in terms of data collection, analysis, expertise, and computer time. Various ad hoc approaches have been used to calculate rough estimates of capacity credit. For example, calculating the capacity factor of the resource over some relevant time period may provide a good general estimate of capacity credit. However, there are few studies that explicitly compare the result of this approach with the more rigorous, traditional approach that is based on sophisticated utility planning models, and none that provide a comparison using alternative numbers of load hours. The purpose of this discussion is to introduce the ELCC and to evaluate simpler methods of calculating capacity credit for an intermittent renewable power plant. These methods should be valuable to a wide audience, including potential investors in, and owners of, renewable power plants, and utilities that are evaluating renewable resources. This discussion follows Milligan and Parsons [9].

Methods for Calculating Capacity Credit

Currently, when a utility evaluates additions to its generator capacity, different energy sources are compared (see Kahn [10] for a discussion of generation project evaluation). Because generator capabilities vary according to fuel type and the method used to produce electricity, it is helpful to use a measure of capacity that can be applied to all types of plants. For example, the capacity value of a 100-MW coal plant might be equivalent to a 75-MW oil plant. A 300-MW (rated) wind plant might provide the same capacity measure as the 100-MW coal plant.

It is important to link this concept of planning capacity credit to 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 (or genco planning), and is the topic addressed here. Operational capacity credit is the capacity value that could be specified in a transaction between utilities. Over the long run, we would expect that the average operational capacity credit would approach the long-term value.

Utility Production and Reliability Modeling

The standard techniques used to evaluate the reliability of power systems and how these techniques are used to measure planning capacity credit are based on Billinton & Allan [11]. Conventional power plants experience unplanned outages, either because of mechanical or other malfunction. Episodes such as this are called forced outages. It is unlikely that conventional generators will experience a forced outage because of fuel shortages. During extended periods of anticipated low loads, generating units can be taken offline for routine maintenance. There is always a non-zero probability that any single generating unit will be on forced outage. Taking all such probabilities from each generator allows us to calculate the probability that enough generator units are on forced outage that the utility will be unable to meet its load. This probability is the loss-of-load probability (LOLP). Most methods of assessing the capacity credit of a wind plant are based on LOLE. Of course, the utility's goal is to keep this probability or expected value as small as possible, given the trade-off between cost-minimization and reliability. A standard rule-of-thumb is to maintain a loss-of-load expectation of 1 day in 10 years.

Effective Load-Carrying Capability

Using the concepts and techniques from reliability theory, 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 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-simulation or reliability model.

The electricity production simulation model calculates the expected loss of load. The usual formulation is

$$LOLE = \sum_{i=1}^N P(C_i < L_i) \quad (1)$$

based on the hourly estimates of LOLP, and the LOLE is the sum of these probabilities, converted to the appropriate time scale. The annual LOLE can be calculated as

where $P()$ denotes the probability function, N is the number of hours in the year, C_i represents the available capacity in hour i , and L_i is the hourly utility load. To calculate the additional reliability that results from adding wind generators, we can write $LOLE'$ for the LOLE after wind capacity is added to the system as

$$LOLE' = \sum_{i=1}^N P[(C_i + W_i) < L_i] \quad (2)$$

where W_i is the wind power output from the wind power plant during hour i . For later sections of this paper, it will be convenient to rewrite equation (4) to allow for wind power output at several locations:

$$LOLE' = \sum_{i=1}^N P\left[\left(C_i + \sum_{j=1}^{N_w} W_{i,j}\right) < L_i\right] \quad (3)$$

where N_w is the number of wind sites in the analysis, j indexes N_w , and $W_{i,j}$ is the wind power output at hour i from site j . The ELCC of the system is the load that can be supplied at a specified level of risk of loss of load.

$$\sum_{i=1}^N P(C_i < L_i) = \sum_{i=1}^N P\left[\left(C_i + \sum_{j=1}^{N_w} W_{i,j}\right) < (L_i + E_i)\right] \quad (4)$$

Calculating the ELCC of the wind plant amounts to finding the values E_i for which this equation says that the increase in available capacity can support E_i more MW of load at the same reliability level that the original load could be supplied with C_i MW of capacity. To determine the annual ELCC, we simply find the value E_p , where p is the hour of the year in which the system peak occurs. Because LOLE is an increasing function of load, given a constant capacity, we can see from equation (4) that increasing values of E_i are associated with declining values of $LOLE$. Unfortunately, it is not possible to analytically solve equation (6) for E_p . The solution for E_p involves running the model for various test values for E_p until the equality in equation (6) is achieved to the desired accuracy.

Although the level of detail of the input data varies between models, hourly electric loads and generator data is required to calculate LOLE. Common outputs from these models include various costs and reliability measures, although cost data are not used to perform system reliability calculations. Some of the models that are used for these calculations are chronological, and others group related hours to calculate a probability distribution that describes the load level, such as that described by Milligan [12]. It is useful to put this discussion in a graphical context. A reliability curve can be plotted by running the reliability model, altering the load, and plotting the resulting points in a graph as in Figure 5. The upward-sloping

lines show 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. Adding a new generator to the model would shift the reliability curve to the right. The level-of-load increase that can be sustained at the same reliability level is the distance between the two 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.

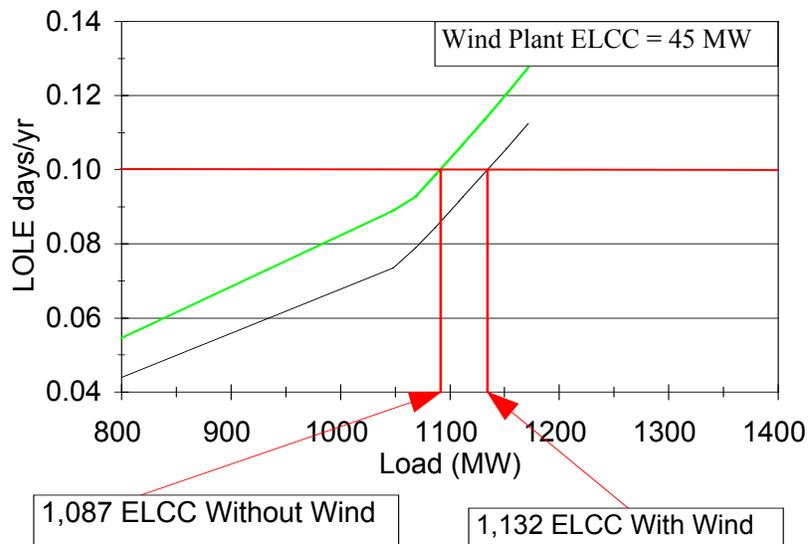


Figure 5. Reliability curves for calculating the ELCC of a wind plant

Application to Intermittent Resources

An important difference between intermittent and conventional generators is the source of the forced outage, as described in the reliability model. From a mechanical standpoint, it is common for wind turbines to have very high availability, often exceeding 0.95. Therefore, the forced outage rate in this case would be 0.05. However, wind turbines can only generate electricity when the fuel (wind) is available. When this is accounted for in the model, the forced outage rate of the wind plant might be in the range of 0.50–0.80. A conventional generator might have an overall forced outage rate of 0.05–0.25, but this rate is likely based on mechanical availability, not fuel availability.

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 [13]), whereas others did not (Bernow, Biewald, Hall, and Singh [14]). Papers by Billinton, Chen, and Ghajar [15], Milligan [12], and Milligan and Graham [16] take related approaches that perform multiple simulations of wind speed, executing the production cost and reliability models over several cases. Ignoring the potential for inter-annual wind speed variation can be perilous, and can result in significantly over- or under-estimating capacity credit.

Simpler Approaches

This discussion has focused so far on standard approaches of measuring power plant capacity credit. Although reliability models provide the most accurate result, they require significant modeling effort. Various ad hoc methods for calculating wind plant capacity credit have been proposed, many of them using the capacity factor over some relevant time period. Related approaches, like those described by Wan, use the median value of the wind plant over a recent history during the utility peak period. This section will provide a benchmark of how well various capacity-factor measures can approximate the ELCC, as calculated by a reliability model.

Application of Alternative Methods

Although we use a wind power plant as an example, these techniques would be equally applicable to other intermittent technologies. The authors calculated hypothetical wind power output based on 13 years of hourly wind-speed data from a site in the Great Plains. This study used the Elfin model with one year of load and generator data from Tri-State Generation and Transmission, Inc. Tri-State is a non-profit generation and transmission co-operative utility, supplying wholesale electric power to 33 distribution co-operatives in Colorado, Wyoming, and western Nebraska (this work occurred prior to Tri-State's acquisition of Plains Generation and Transmission). 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 (WAPA) and Basin Electric Power Co-operative (Basin).

Capacity credit calculations were carried out for each of the 13 years of data using the standard ELCC approach. We then used three alternative methods to calculate wind plant capacity credit. None of these alternative methods require using either a production cost or reliability model or generating data from conventional generating sources, although two of these methods do require a single reliability model run. The first method calculates the capacity factor (defined as the ratio of the average output to the total output) for the hours during the utility system peak. This calculation was carried out for the top 30% of hourly peak loads.

The second method also calculates the capacity factor, but uses hours in which the risk of not meeting the load is highest. These hours are determined by running the reliability model with no wind generation and calculating the LOLP for each hour of the year. The hours selected by this method generally correspond to the hours of the highest loads, but differ to the extent that conventional resource availability is not constant throughout the year. The wind-power output for these hours is then used to calculate the wind plant capacity factor as an estimate of the wind plant's capacity credit. These calculations are carried out for up to 30% of the yearly hours.

The third and final method used the same hours as method two. This final method uses normalized LOLP values as weights for the average capacity factor. This allows the method to recognize those hours in which LOLP is more severe and weight them accordingly. The capacity factors are then calculated in the same way as those in the other approaches.

The wind data used came from a regional air quality monitoring program (RAMP) site in the high plains, scaled to match average wind speeds from the Altamont Pass region in California. The series spans the years 1980 to 1992. The wind data collected from this site show a wide inter-annual variation. The highest simulated energy capture occurred in 1988. Values for other years range from about 72% to about 92% of the 1988 simulated energy. A caveat to our approach is that several hundred miles separate the utility and

the wind site, and the utility data is for a single year. This implies that any correlation between loads and wind is not fully captured. This is a limitation that is likely in practice, because of the relative scarcity of long-term wind data.

Results

Each calculation method was applied to the full 13-year data set. For each year of wind data and each method, the wind plant capacity factor was calculated for the top 1%, 2%, . . . , 30% of hours that were selected by the algorithms described above. Figures 6–9 illustrate some of the variation that was found. Figure 6 shows the 1980 capacity factors for the three methods. The “Load” series was created by taking the top system loads, the “LOLP” series by taking the hours for the maximum value of hourly LOLP, and the “Weighted” series uses normalized LOLP as the weights for the capacity factors. The ELCC value of 31.3 was calculated using the reliability model and is the benchmark value used in the study. Figures 7, 8, and 9 are constructed in the same way for the years 1982, 1984, and 1990, respectively.

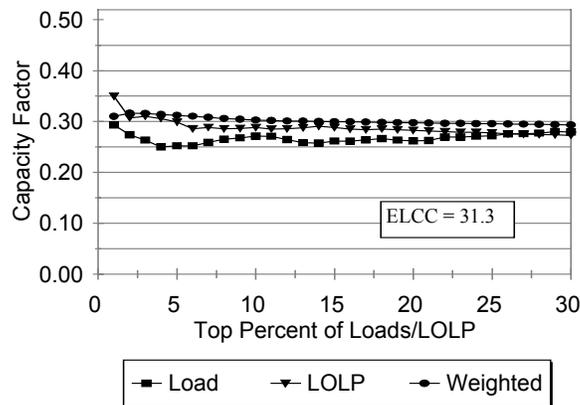


Figure 6. Capacity factor for top hours, various methods, 1980

Points to the right of the 10% level of Figure 8 for all three methods appear to slightly underestimate the capacity credit, as calculated by ELCC. The weighted approach appears to be more accurate at low-load percentages, but all three methods provide similar capacity credit estimates as we approach the 30% load level.

The results for 1982 are depicted in Figure 7 and are quite different. There is clear disagreement between method 1 and methods 2 and 3. However, this result should come as no surprise. The ELCC measure is based on system reliability, which is a function of both load level and available capacity. Utility dispatchers will attempt to provide as much capacity as possible during system peak. Non-peak periods may have higher LOLP values, and, thus, lower reliability, if lesser capacity is available. In Figure 7, there is quite a close agreement between the two methods that take LOLP into account, and these methods provide approximations that are very close to the ELCC value over a wide range of top load hours.

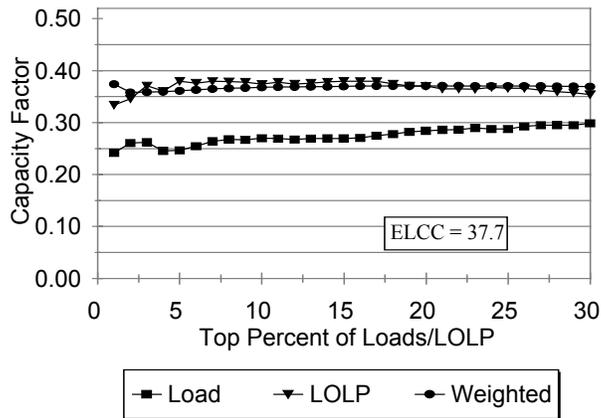


Figure 7. Capacity factor for top hours, various methods, 1982

Figure 8 shows what can happen when wind generation during the utility’s extreme peak period is not typical of the wind generation over lesser peak periods. Choosing hours to evaluate based solely on load level severely underestimates the ELCC, although when a larger number of hours are included in the calculation, this error quickly tails off. However, both of the LOLP-based methods substantially overestimate the capacity credit. Method 2 converges much more quickly than does method 3. When the top 30% of hours are used, all methods are quite close to the ELCC value.

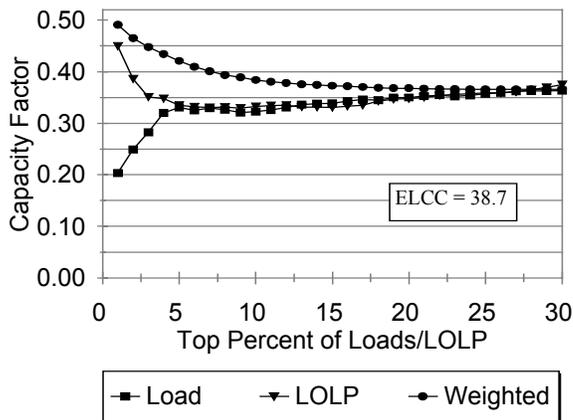


Figure 8. Capacity factor for top hours, various methods, 1984

Figure 9 is the final graph of this series. In this case, method 1 severely overestimates capacity credit when a small number of hours are used, whereas methods 2 and 3 are consistently below the ELCC value. These series do not converge very well, either to each other or to the ELCC value. Each of these methods underestimates capacity credit for 1980.

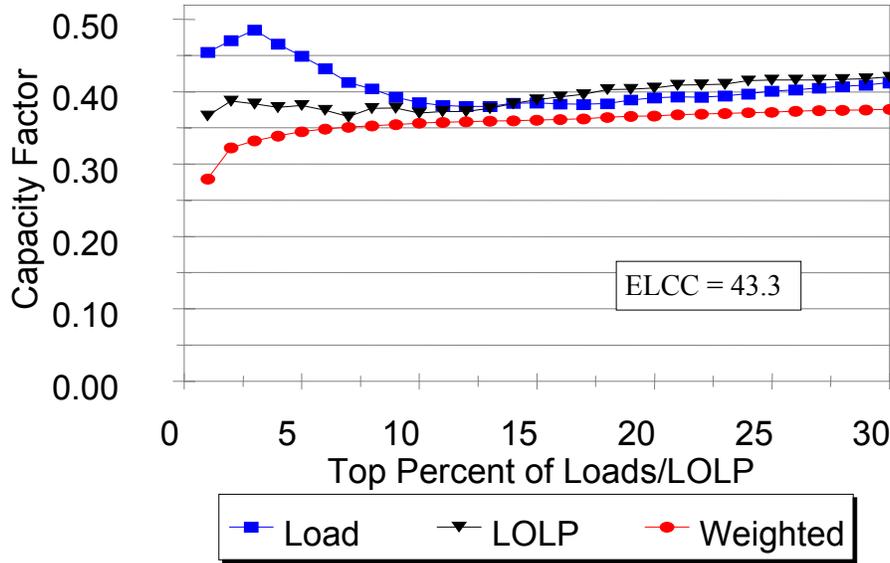


Figure 9. Capacity factor for top hours, various methods, 1990

We can draw several conclusions from the results. First, although capacity factor might be useful as an approximation to capacity credit, it appears to consistently *underestimate* the ELCC value. Second, the accuracy of these capacity factor methods is very sensitive to both the number of hours used and the method used to select the hours. Third, wind power plants contribute to overall system reliability during non-peak hours.

Comparison of Methods

Although it is apparent that these capacity factor methods are not as accurate as ELCC methods for calculating capacity credit, Milligan and Parsons performed some simple calculations to rank these methods. For each of the 13 years of wind data and for each method, we calculated a root-mean-square error (RMSE) statistic:

$$RMSE = \frac{1}{n} \sqrt{\sum_1^n (x_a - x_f)^2} \quad (5)$$

where n = the number of hours used in the approximation, x_a is the actual capacity credit (as calculated using the ELCC approach), and x_f is the estimated capacity credit by method 1, 2, or 3. The RMSE was only calculated on the results based on the 30% of load hours. Figure 10 represents a simple scoring of the

results—how many years did each method produce the lowest RMSE—alongside the average RMSE over the 13-year period. Although method 1 is tied with method 2 for the number of years with the lowest error, its average RMSE exceeds that of method 2. In spite of delivering impressive performance in 1982, method 3 has the highest composite error.

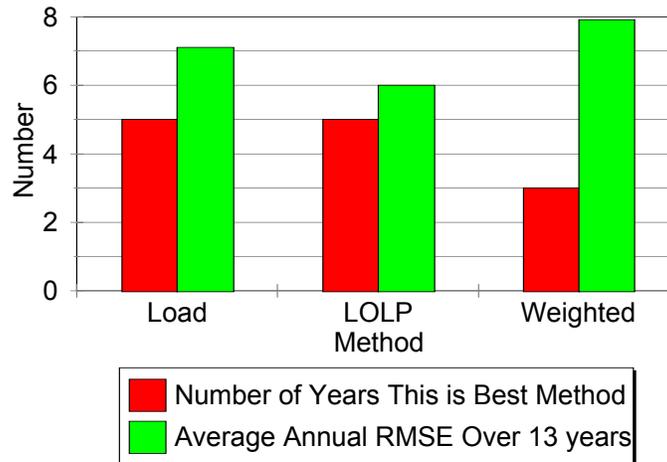


Figure 10. Overall comparison of methods

In our judgement, the Load method provides a reasonable trade-off between accuracy and effort, either early in project assessment or if it is not possible to calculate the ELCC. In other situations, it might be possible to obtain LOLP output from a reliability model, but not possible to perform multiple runs to calculate ELCC. Here, the LOLP method performs better with a modest effort. Although the weighted method does not appear to perform well here, this is likely because of the lack of correlation between load and wind, an artifact of the data sets available for this study.

Summary

To fully measure the long-term capacity credit of an intermittent power plant, Milligan and Parsons recommend using the standard ELCC measure and a full complement of reliability model runs. However, there may be cases where this is impossible. If hourly LOLP values can be obtained from a reliability model run but multiple model runs can't be performed, they recommend using the LOLP method. If no reliability measures are available, the Load method appears to provide a reasonable approximation. There is clearly a trade-off between accuracy and effort. As the restructuring of the utility industry moves forward, it will be important to continue to examine different calculation methods that can be easily applied by various players in the industry.

So far, we have seen the relationship between wind forecasting, unit commitment, and operational capacity credit. An accurate forecast can provide a benefit to the utility or generating company by allowing optimal, least-cost-unit commitment. In the context of the restructured market, an accurate wind forecast can also be useful to the wind plant operator because it helps specify the timing of wind capacity that can be contracted for in a day-ahead market. With an accurate wind forecast, wind capacity can be sold in the market.

Measuring capacity credit with ELCC is a relatively common approach. The Milligan and Parsons work demonstrated the differences obtained when comparing capacity factor and capacity credit based on ELCC. However, none of the work we have seen so far has solved the problem inherent in both Elfin and P+: how can we specify hourly wind power data *and* obtain an estimate of the wind plant reliability that is caused by variation in wind speed? That question is addressed in the following section (NREL has developed a probabilistic method for incorporating wind power plants into the system-reliability calculation. This will be published in the near future).

Sequential Monte Carlo

Intermittence and the high variability of wind make it difficult for models to adequately measure capacity credit. Capacity credit results depend heavily on what happens during the utility's peak hours. Because wind speed can vary significantly from year to year and from hour to hour, capacity credit estimates that are based on a single year (or less) of data and modeled without taking this variation into account may not be credible.

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 multiple-scenario or 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 [17]. Milligan [12] illustrates a Monte Carlo method that is external to the load-duration curve production cost model. This approach creates a set of many wind power series, each of which can be run in the production cost or reliability model, and is similar to a technique proposed by Billinton and Chen [15]. Milligan [18] applies this approach to the RAMP wind site, and compares the capacity credit results obtained with the external Monte Carlo method with results using the actual wind-speed data. Milligan and Graham [16] and [19] extend the basic framework, using the Elfin and P+ models, and introduce a reduction technique to help minimize the significant model run-time that is required for the full simulation set.

This section summarizes these various modeling approaches and results. After a short discussion of the external Monte Carlo method, we begin by illustrating the method and show how multiple wind-speed and wind-power simulations compare with the 13-year RAMP data set, following Milligan [18]. The discussion will then turn to applying the method to the two production-cost models, and finally, to the reduction method proposed by Milligan and Graham.

Many production-simulation and reliability models have a Monte Carlo option that allows sampling from the generators availability probability distributions. This approach is used when one would like a better estimate of the range of possible outcomes than can be provided by the usual convolution approach. Another advantage of the Monte Carlo method is that it provides estimates of various probability distributions, such as system reliability and system costs. The P+ model also has a branching option, which combines the more efficient convolution approach with the more precise Monte Carlo method. The branching technique performs the usual convolution on all but one generator. This generator's state will be sampled repeatedly via Monte Carlo, holding all other generator values to the expected values from the convolution. This allows the analyst to focus on the effects of a particular generator, without paying the full price of heavy execution time that is exacted by full Monte Carlo.

This approach appears to be ideal for modeling wind power plants. Unfortunately, the Monte Carlo simulation procedures generally sample from a very simple probability distribution that is not appropriate for wind power plants. This leads us to consider separating the probabilistic sampling from the production-cost model. The method involves repeatedly creating synthetic wind-speed data, which can easily be used to calculate hourly wind power output. One can obtain a sequence of such data sets, and then run a series of production model simulations, capturing the results of these runs and summarizing them in a convenient form. The Monte Carlo is used to create the synthetic wind series, and the production-cost or reliability

model can be applied to each. This is sometimes called “sequential Monte Carlo,” to differentiate it from the Monte Carlo logic that is often found in the models themselves.

This solves the problem that is inherent in the somewhat unappealing trade-off between maintaining hourly wind power data and providing a probabilistic measure of the event that the wind power may vary from a fixed value. However, the cost of this solution is increased complexity and model runtime.

A Markov Model for Wind-Speed Variations

In this section, we examine the Markov model, which captures the time-varying properties of wind-speed. There are several variations of the Markov model. In other work, I have tested some of these methods, and found that the choice of method should be based on the characteristics of the wind site [20]. Parsimony suggests choosing the simplest model that performs adequately. Sites such as the HP site cannot be modeled any more effectively with more advanced Markov models than with simpler ones.

A stochastic process is defined as a Markov chain if [21]:

$$(P(X_{t+1} = j \mid X_0 = k_0, X_1 = k_1, \dots, X_{t-1} = k_{t-1}, X_t = i) = P(X_{t+1} = j \mid X_t = i)) \quad \forall t = 0, 1, \dots \quad (6)$$

and all sequences $i, j, k_0, k_1, \dots, k_{t-1}$. This definition says that the conditional probability of any event, given any past event and the present state $X_t = i$, depends only on the present state of the process. In the context of wind speed, this implies that the probability of any wind speed for hour t depends only on the wind speed in hour $t-1$.

Some wind regimes, among which are several sites in California, have strong diurnal components. Other sites, such as the HP site, have wind that is driven primarily by the passage of frontal systems. In this latter case, one would expect a very weak diurnal pattern, and the resulting Markov assumption can be written as follows:

$$P(X_{t+1} = j \mid X_t = i) = P(X_1 = j \mid X_0 = i) \quad \forall t = 0, 1, \dots \quad (7)$$

Equation (7) says that the probability distribution that describes transition from one state to the next is independent of time. Such a process is called a stationary Markov process. We have tested both stationary and non-stationary Markov processes.

Of particular interest is the obvious restriction that the Markov assumption places on the memory of the process. Other methods, such as the auto-regressive integrated moving average (ARIMA) approach originally put forward by Box and Jenkins [22], allow for generalizing the influence of past events on the current event. The first-order autoregressive model, the AR(1), is a special case of a stationary Markov model. The Box-Jenkins approach involves analyzing the data series and its autocorrelation function to determine the order of the autoregressive process, which can be greater than 1 (see [15]). We tested a second-order Markov process, which can be described as:

$$P(X_{t+1} = j \mid X_0 = k_0, X_1 = k_1, \dots, X_{t-1} = k_{t-1}, X_t = i) = P(X_{t+1} = j \mid X_t = i, X_{t-1} = m) \quad \forall t = 0, 1, \dots \quad (8)$$

This approach assumes that the probability of the state of the process depends on the previous two states. This allows the model to implicitly take the trend in wind speed into effect, a feature that is missing in the basic Markov model.

It is well known that seasonal variations in wind speed can be significant. For this work, Milligan and Graham applied each method independently for each month of the year. Although the Gregorian calendar is somewhat artificial with respect to meteorology, it does allow us to capture changes in wind speed that are a function of the time of year.

Each of the Markov variant methods calculates a set of state-transition matrices. Each matrix is calculated for a single month using all available wind-speed data. The state-transition matrix is a series of probability distributions. One distribution is calculated for each integral wind-speed value (measured in meters per second). Therefore, for each wind-speed value, we have the probabilities of changing to all possible wind-speed states. After calculating the state-transition matrices, the simulated wind-speed series are constructed by sampling from the appropriate probability distribution, drawing a random wind-speed value for each hour.

Most of the work described here uses wind data from the Great Plains. The Markov model used in these cases is the basic first-order, stationary process, and is selected based on an RMSE goodness-of-fit criterion.

Figure 11 illustrates a state-transition matrix for July. During the simulation process, twelve such matrices are created—one for each month of the year. The next step is to sample from the state transition matrix, thereby creating an artificial wind-speed series based on the original series.

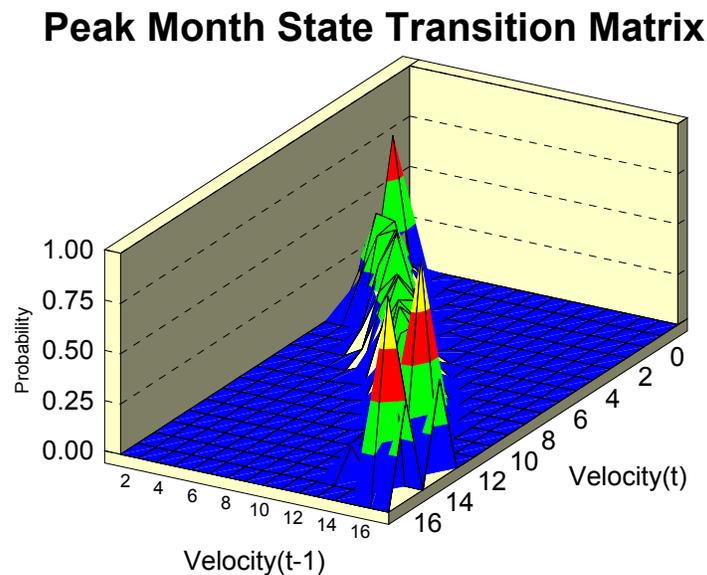


Figure 11. Peak month state transition matrix

This process normally requires several wind-speed series. Based on work we have done at NREL, at least 1,000 series, in our judgement, will provide an adequate representation of wind-speed variation. Billinton, Chen, and Ghajar report that, using an AR(2) Box Jenkins model, their model converges to a specified tolerance in about 6,000 iterations.

Capacity Credit Modeling with Elfin: Comparison of the Markov Model to Actual Wind-Speed Data

Milligan [18] applied the Markov model to the 13-year RAMP data set. Using load and generator data from Tri-State Generation and Transmission Association, Inc., I calculated energy output from a hypothetical 100-MW wind power plant, and used this as input to Elfin. Figure 12 shows the annual energy and ELCC, each based on the percentage of 1988, the year in which maximum energy and capacity credit occurred. As can be seen from the graph, there is a high correlation between wind energy output and ELCC; however, there are several times that these series move in opposite directions.

Capacity Credit and Energy for 13 Yrs

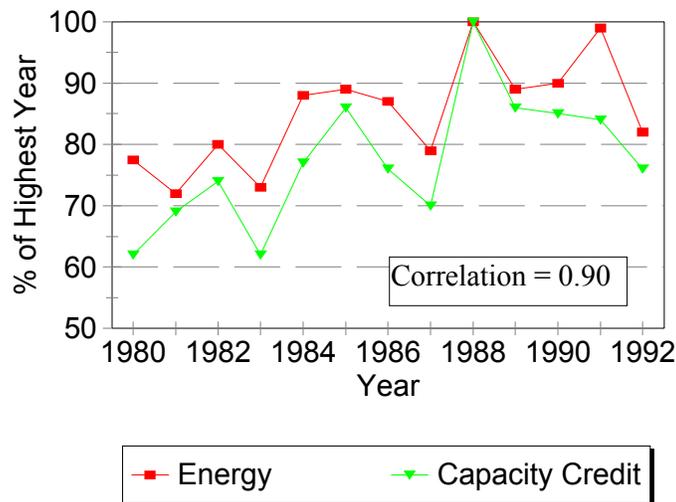


Figure 12. Capacity credit and energy for 13 years

Representative years are selected from the 13-year data set, on which a Markov simulator tool is applied. This allows us to analyze plausible wind-speed variations and their effect on capacity credit. Based on the actual data, 1980, 1986, and 1990 were selected to represent low-wind, average-wind, and high-wind years. Using these three years, separate state-transition matrices were calculated, and a set of 100 simulations was done using each one.

Figure 13 shows the combined results. The first column in the graph shows the ELCC values based on the actual wind data. The mean ELCC over the 13-year period is 39.1% of rated capacity, with maximum and minimum values of 50.5% and 31.3%, respectively. The range of values (50.5–31.3) is 19.2, and the coefficient of variation (COV) is 13.8%. As one might expect, the high-year simulations yielded a higher maximum ELCC (52.5%) with a correspondingly higher mean (42.3%). However, the minimum is slightly lower than the mean of the actual data (30.5%), and the COV is 8.6%, which is markedly lower than the actual series COV. The mean-year simulation is a closer match to the actual year. The mean is slightly higher, and the range is somewhat higher. However, the COV is only 10.1%, as compared to 13.8% for the actual data. Results for the low-year simulation exhibit similar tendencies. The mean, maximum, and minimum are lower than actual. This results in a higher range and a lower COV. These results suggest that the Markov approach provides a reasonable look at how variation might have an impact on ELCC, but that it generates a small number of more extreme events (with a higher range than in the actual data) while underestimating overall variation (a lower COV as compared to the actual data).

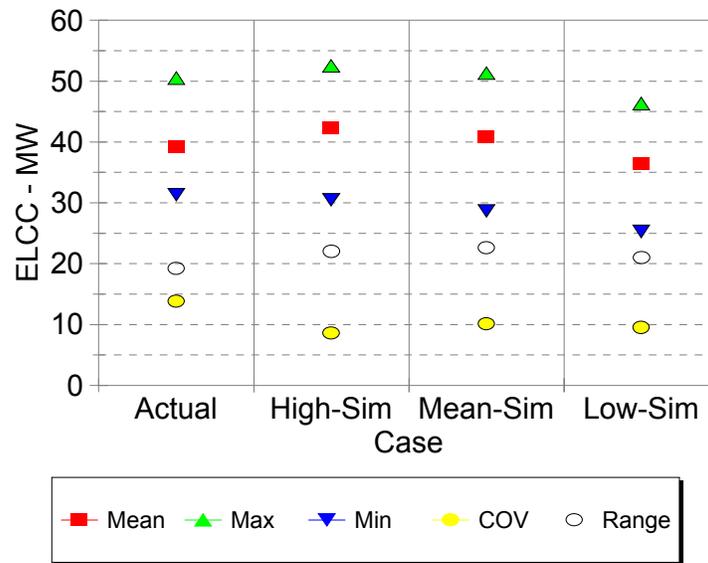


Figure 13. Comparison between simulated and actual ELCC

The energy results can be seen in Figure 14. The energy variations that are simulated by the Markov process clearly underestimate the natural variation found in the actual data set, as evidenced by markedly lower ranges and COVs. The energy COV values from the simulation are nearly one-half of the actual COV. Clearly, the Markov approach underestimates energy variations.

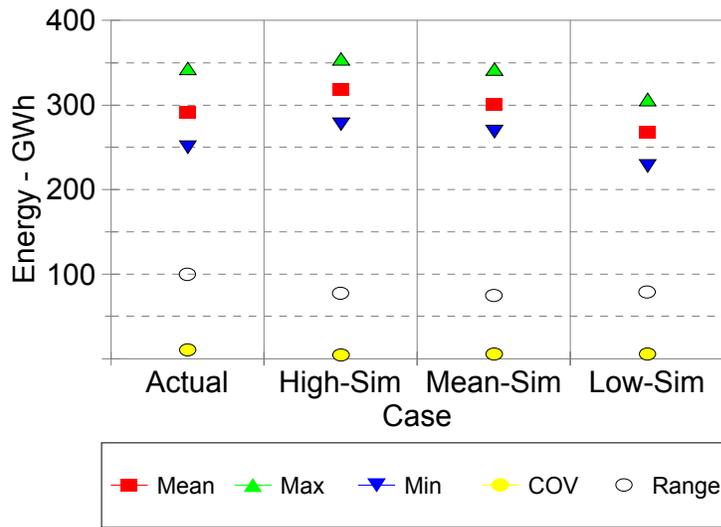


Figure 14. Comparison between simulated and actual energy

Although not ideal, the Markov model and its integration into sequential Monte Carlo simulations provide significant insight into expected variations that might arise because of annual wind-speed differences. As more wind data becomes available in the public domain, additional research can provide better models and validation. If only a single year of wind data is available for analyzing a particular project, an attempt could be made to compare this data with other long-term data records to determine how the measured year compares to a longer period. Once this has been established, the Markov method could be used to calculate the estimated variations.

There are several caveats to this work. First, these results may not be robust to other wind sites. The data used here does not have a significant diurnal component, as the wind is primarily caused by synoptic-scale weather events. This data was obtained from an air-quality monitoring site that would not be used for a wind power plant. It may also be the case that an even longer data set is necessary to adequately cover the range of expected variation. Capacity credit, however measured, is sensitive to the correlation between the utility load and wind resource and to the specific utility resource mix during peak and off-peak periods. ELCC is also sensitive to the risk-target chosen by the utility management, and whether LOLE or ENS is used as the reliability measure.

Capacity Credit Assessment Using Sequential Monte Carlo: P+ and Elfin

Milligan and Graham illustrate how the Markov-Sequential Monte Carlo technique can be applied to both load duration curve (LDC) and chronological production-simulation/reliability models. Part of this work focused on a reduction technique that selects among the many wind-speed and wind-power time-series that are generated from the Sequential Monte Carlo method. To distinguish the reduction method, the Enumerated Probabilistic Approach (EPA) denotes the full Sequential Monte Carlo method, whereas the Reduced Enumerated Probabilistic Approach (REPA) designates the reduction technique. We used the Elfin model to establish a base case of 1,000 EPA simulations, then used both Elfin and P+ to perform a

set of REPA simulations using a subset of the 1000 simulations. The paper contains some selected reliability and cost results, and compares the EPA and REPA methods.

The utility data used in this study is from Tri-State Generation and Transmission Association, Inc. To provide a plausible analysis of wind plant reliability and ELCC, a Markov wind-speed simulation tool was applied to a single year of wind data from the Nebraska Energy Office. We chose the Imperial, Nebraska, site because of its proximity to Tri-State's service territory.

This analysis focuses on October 1995, a month in which there appears to be significant variability in the wind resource. To satisfy both models' requirement for 168-hour weeks, each model was run for six full weeks before obtaining 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 the load by 90 MW to account for a time-varying purchase from Basin. The net peak load was 1,350 MW. The maximum hydro purchase from WAPA was 400 MW, with 1,152 MW of base and intermediate generation and 120 MW of peaking capacity. Milligan and Graham modeled a hypothetical 100-MW (nameplate) wind plant.

The wind-speed state transition matrix for October appears in Figure 15. This graph shows the probability of occurrence of each wind speed at time t as a function of velocity at time $t-1$. 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. The method could be appropriately applied to other time scales. Once the multiple wind-speed realizations have been simulated, one can calculate the hourly wind power output from a hypothetical wind plant for each realization. We then performed the analysis of all wind speed realizations (EPA) or of the cases selected for the REPA analysis.

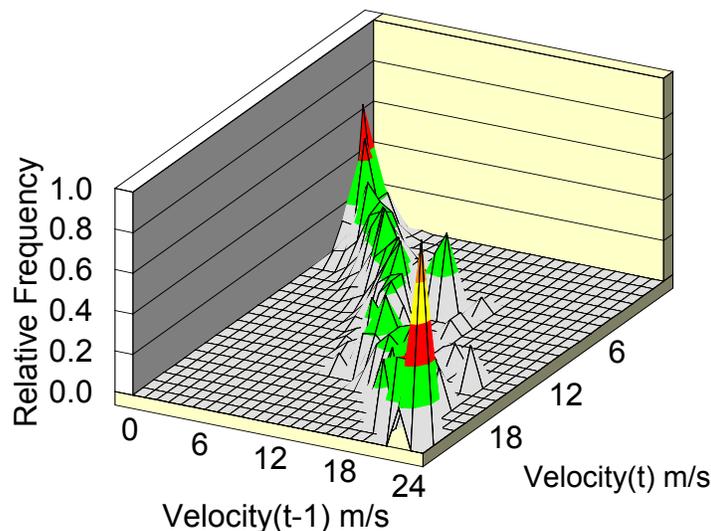


Figure 15. Wind-speed state transition matrix for October

For the EPA runs, each hypothetical wind power realization is input to the Elfin model, which is executed for each one. From this process, one can obtain the ELCC of each realization, which can then be summarized for further analysis.

The REPA method is an attempt to reduce the number of reliability model executions with a minimal reduction of accuracy. The approach is to group the 1,000 wind-power series based on energy output during the utility system peak. The data are then grouped and assigned weights corresponding to the wind energy frequency distribution. Representative wind-power realizations are then selected from each of the groups, and the model is run once for each selected case. Weighted averages are then computed for the outputs of interest. We performed this analysis with both the Elfin and P+ models.

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.

The process of selecting the various wind realizations for the REPA involves some judgement. The 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 five bin sizes corresponding to the mean and 1 and 2 standard deviations resulted in a relatively poor representation of the variation found in the ungrouped cases. Faced with a trade-off between execution time and accuracy, Milligan and Graham 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. However, 11 bins were used, each having a width of one-half the standard deviation of wind energy produced during the utility's peak. Basing the bin selection on all cases within 2.5 standard deviations of the mean energy allowed us to retain the variation in the full EPA data sets, but with a significant reduction in the number of wind-speed realizations. Figure 16 illustrates the distribution of the wind energy during the utility's peak period in October. The authors believe that the choice of bin width and number did a good job of capturing the variation in the full EPA results.

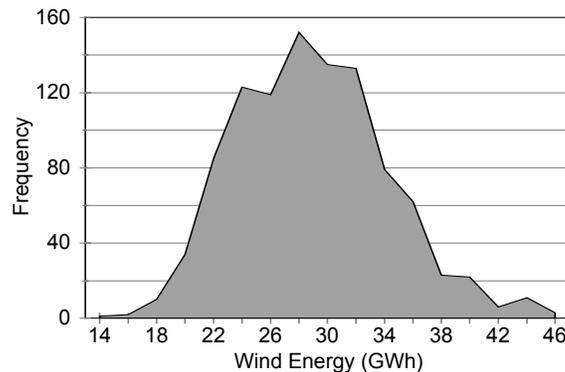


Figure 16. Wind energy distribution for multiple data sets, October

Figure 17 shows the bins and resulting weights that were selected for the analysis. This grouping retains the shape of the original distribution, while adequately representing the variation in the data.

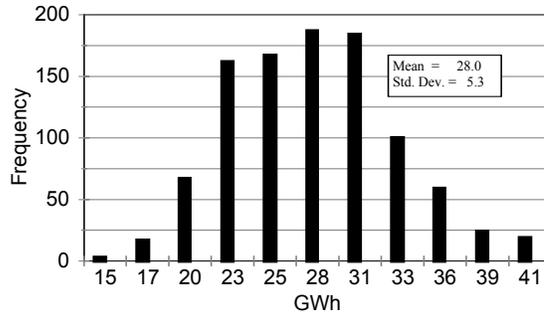


Figure 17. Wind energy distribution, binned by standard deviation

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 using Expected Unserved Energy (EUE) and Loss-of-Load Hours (LOLH). It is clear from the table that the EUE reliability measures are in closer agreement than the LOLH measures. On a percentage basis, the EUE difference is about 4%, whereas the LOLH difference is about 5%. In our judgement, the EUE is likely to be more accurately estimated than LOLH, as measured by the two models, and this is what is used as the basis for the ELCC calculations.

Model	Expected Unserved Energy (GWh)	Loss-of-load Hours
Elfin	6.6	46.1
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 the equivalent. The ELCC values that were calculated are based on those calculated in the base cases illustrated in the table.

To take full advantage of Monte Carlo simulations such as the 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, it was not possible to specify convergence criteria; only the number of runs to perform. See Marnay and Strauss [17] for further discussion.

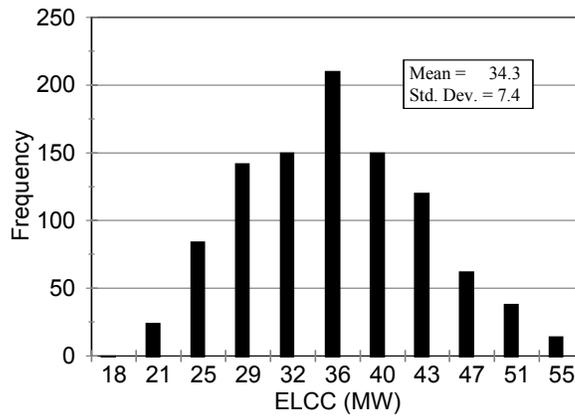


Figure 18. Capacity credit distribution of EPA runs

For the 1,000 simulations of the EPA, Figure 18 illustrates the range of ELCC values as a percentage of installed wind capacity. Although these results appear to be consistent with the annual results reported in Milligan [18], here there is a larger variance of capacity credit because of the larger variation in wind plant output over the month than would occur over a year. Each bin for the figure represents a width of $s/2$, where s is the sample standard deviation. Our data indicates that all but about 2% of the values are within two standard deviations of the mean.

From the 1,000 cases run for the EPA method, cases were identified that most closely matched the mean ELCC, and the mean plus or minus $s/2$, s , $3s/2$, $2s$, $5s/2$ of ELCC. This resulted in 11 of the wind-power realizations from the full data set, which were used to perform the REPA analysis.

The results of the ELCC calculations are presented in Figure 19. 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. The monthly capacity factor is also shown for comparison, and is virtually the same as the full EPA ELCC value. The standard deviation of the EPA and REPA differ somewhat, as shown in Table 2.

Statistic	P+ REPA	Elfin REPA	All 1,000 Cases: Elfin EPA
Mean	33.4	31.9	34.7
Std. Dev.	11.8	9.9	7.4

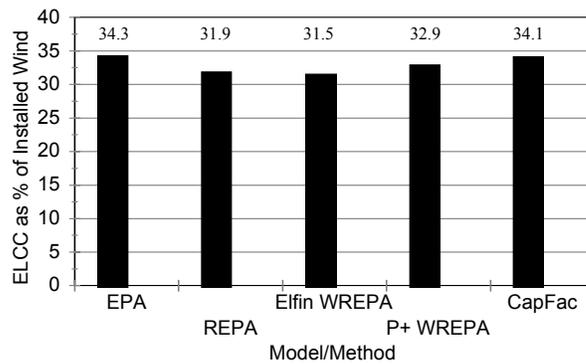


Figure 19. Capacity credit results from various methods

The table indicates that the representative cases overestimate the variation of the larger sample, as indicated by the standard deviations. Conversely, the reduced-sample mean value appears to underestimate that of the larger sample. Although our REPA process exhibits a higher standard deviation than the full EPA runs, this still likely underestimates the monthly variation that could be expected from year to year.

Figure 20 shows the ELCC results for all of the REPA cases for both models. The case numbers in the diagram indicate increasing wind energy levels that correspond to intervals chosen for the REPA analysis. Although one would generally expect that higher wind energy levels would result in higher ELCC estimates, the graph shows that this is not entirely true. ELCC is a function of the generating system's reliability level, which is in turn a function of load and available capacity. As described in Milligan and Parsons, higher wind energy levels during a specific period will not always increase the ELCC of the wind plant. This can also be observed in Figure 21, which plots energy vs. ELCC of the wind plant. The upward trend shows the positive correlation between higher energy and ELCC values, but there are cases in which higher energy will not correspond to higher capacity credit. The diagram also shows a difference in the capacity credit from the two models. This is likely caused by the difference in commitment algorithms used by chronological and LDC models, as discussed in an earlier section of this paper.

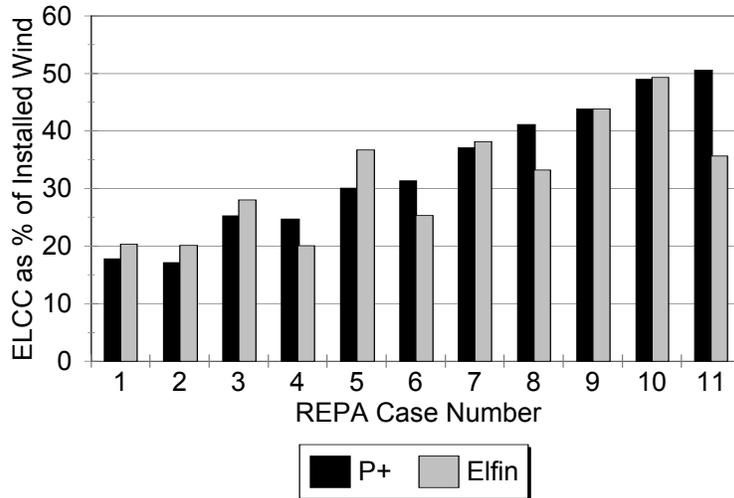


Figure 20. Comparison of the REPA cases by model

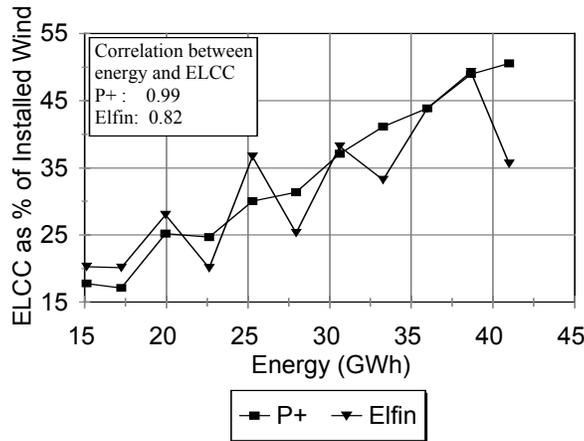


Figure 21. Correlation between wind energy and ELCC

It is also important not to lose sight of the objective, which is to develop a computationally efficient way to provide plausible estimates of wind plant output and *variation* in output. To that end, the REPA method has accomplished that goal. Further work should be done to explore the impact of the utility's peak period and overall reliability level on these results.

Summary

The REPA appears to be a computationally efficient way to examine the impact of possible variations in wind plant output. Instead of implementing or modifying a Monte Carlo routine that is embedded in a production-cost or reliability model, the Milligan and Graham method provides the model with a small number of wind power data sets. The model can then be run for each of the (reduced) 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 reproducible method that could be converted into a computer algorithm. In our judgement, the bin sizes and selected ranges are reasonable. Although the REPA is not as accurate as the EPA, it does capture the variation in the wind resource.

Modeling Production-Cost Variations with Sequential Monte Carlo

Although it is not part of a capacity-credit assessment, it is useful to view other results from the Sequential Monte Carlo model runs. As a chronological model, P+ can produce hourly results for each day and week. Figure 22 illustrates the weighted hourly change in generation for a 1-week period. This week was chosen to illustrate some substantial variation in hourly wind power output. Other weekly model outputs are similar. Using the REPA approach and bin selection process results in 11 similar graphs, one for each bin, each representing various plausible scenarios. Generation planners and analysts then have a range of such outputs on which to base their decisions. Figure 23 shows the reduction in conventional generation for a typical day. The solid line shows the weighted average of the reduced data set, and the other lines show selected results from two of the bins.

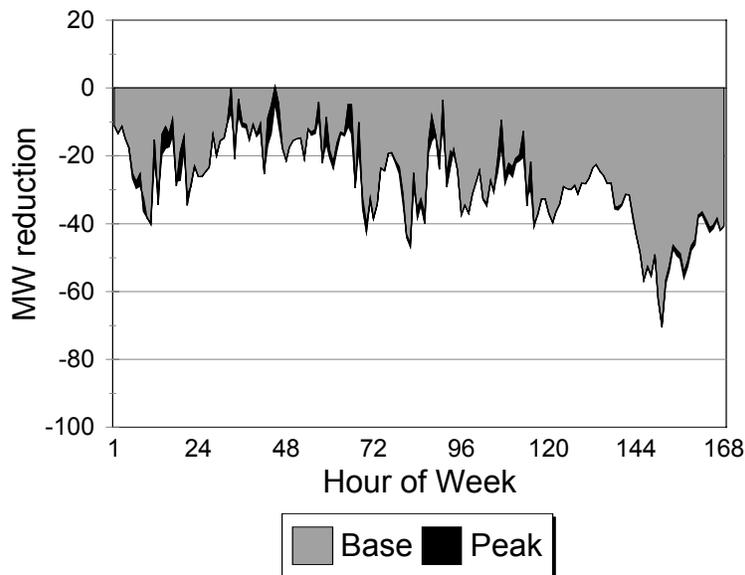


Figure 22. Generation reduction, week of October 9

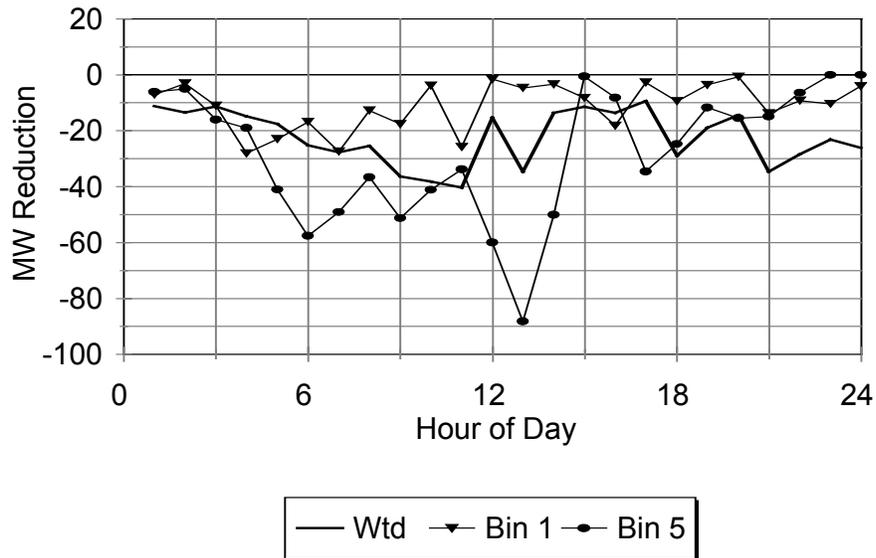


Figure 23. Generation reduction, October 9

Milligan and Graham [19] perform a similar analysis, but with only five bins (they prefer the larger number of bins, but the use of five bins will help clarify these examples). Each bin width corresponds to the standard deviation of the wind energy output from the EPA cases, and is therefore twice as wide as those discussed in Milligan and Graham [19]. The reduction technique provides a very good summary of a wide variety of model outputs, depending on the analyst's interest. We have chosen to include graphs that show daily and weekly results of the scenarios represented by each of the five bins. Figure 24 shows changes in the base load and intermediate units in each scenario. Figure 25 summarizes this information, showing the mean value for the five REPA cases. The weighted values are calculated with the probabilities associated with each of the bins. The next two figures, 26 and 27, show similar information for the peaking units.

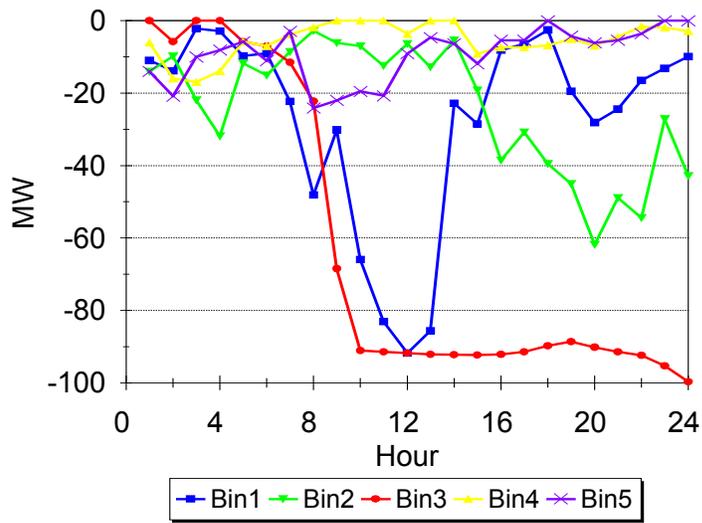


Figure 24. Change in base and intermediate generation from no-wind case, REPA cases, October 14

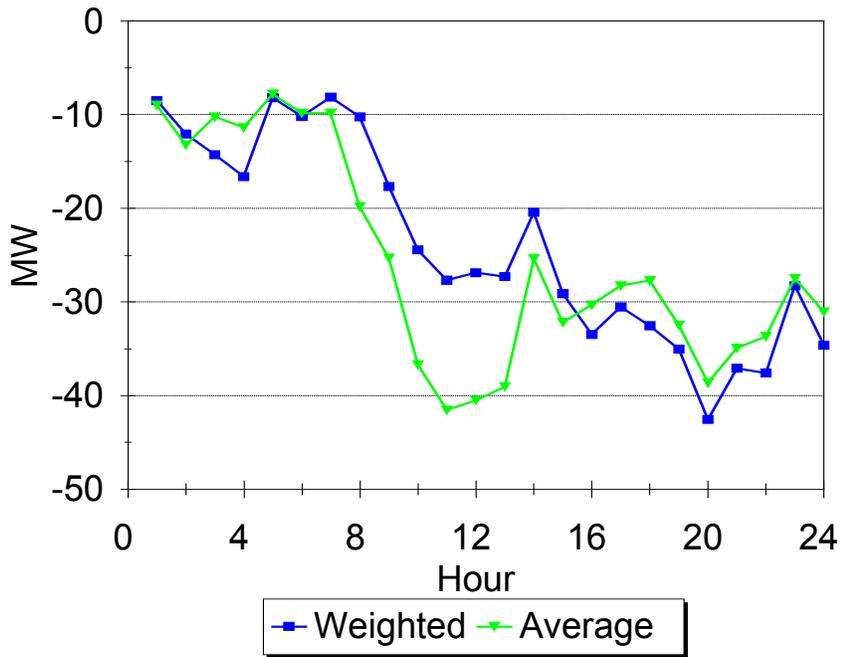


Figure 25. Change in base and intermediate generation, mean REPA and weighted REPA cases, October 14

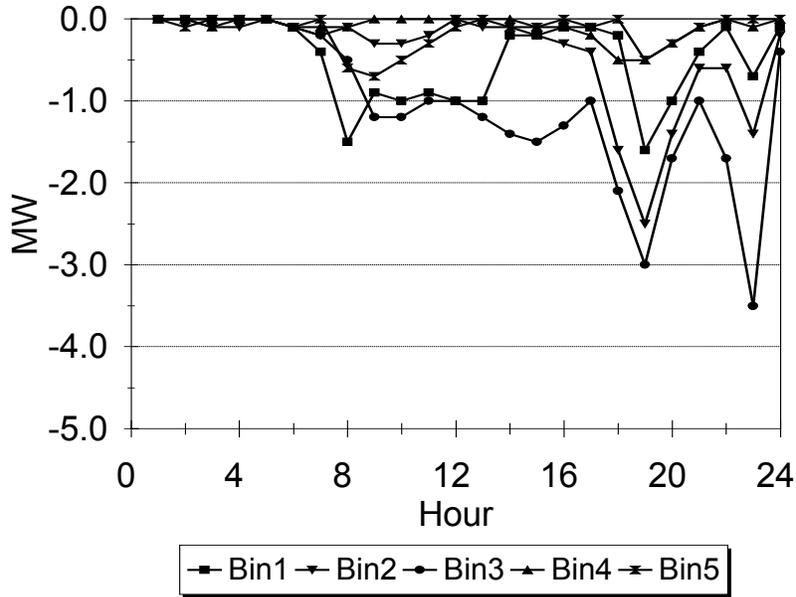


Figure 26. Change in peaking generation, REPA cases, October 14

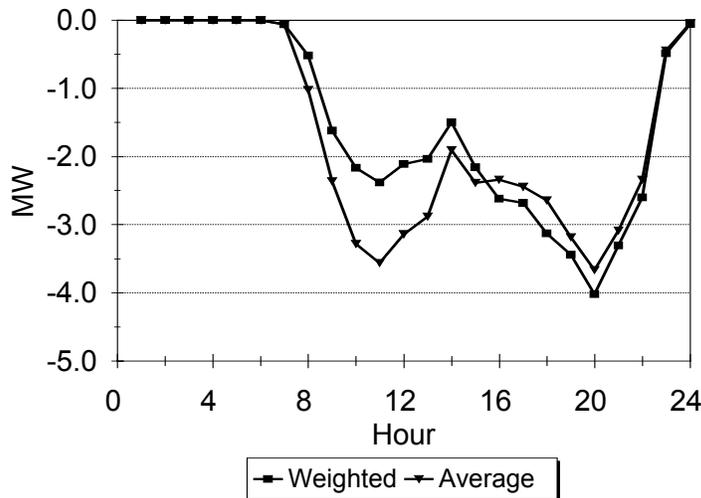


Figure 27. Change in peaking generation, mean REPA and weighted REPA cases, October 14

Weekly results can also be shown so that one can obtain an idea of how the wind power variations influence the generating system over a longer time period. As illustrations, the authors have chosen a graph that shows the change in weekly generation, Figure 28. This diagram shows the differences between the five representative cases and the weighted average results of the binned cases. Figure 29 shows a similar range of reliability values based on EENS.

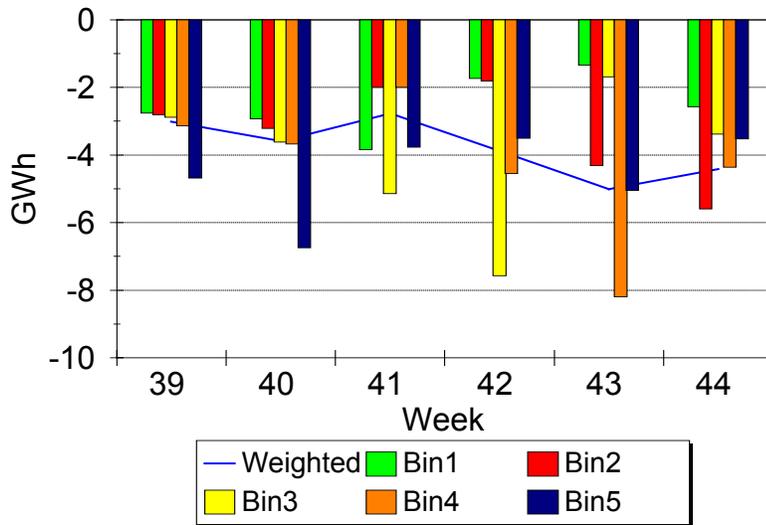


Figure 28. Weekly change in base and intermediate generation for 5 REPA cases

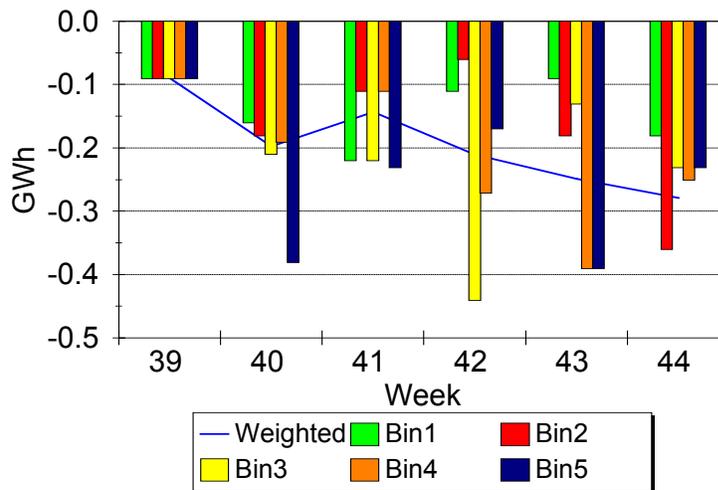


Figure 29. Weekly change in EENS for 5 REPA cases

Summary

Calculating the capacity credit of a wind-power plant can be a complicated process, requiring extensive data and sophisticated models. The Milligan and Parsons results suggest that calculating wind plant capacity factor for the top 10% of loads over the year can provide a reasonable estimate of ELCC, although in some cases there may be a significant loss of accuracy. Even a more sophisticated method that calculates ELCC may be suspect if there is significant inter-annual variation in wind-speed and timing through the year. The Sequential Monte Carlo method provides a very powerful way to evaluate the risk that intermittent resources impose on the generating system. As wind systems become more prevalent in the emerging restructured utility industry, methods to help understand plausible variations of wind power plants and their effects will prove important. The reduction technique does a very good job in screening a few cases for further analysis, and reasonably represents the variation that is likely in wind power over the year.

Modeling Geographically Disperse Wind Power Plants

There are several wind plants in various stages of planning or development in the United States. Although some of these are small-scale demonstration projects, significant wind capacity has been developed in Minnesota, Wyoming and Iowa. As these and other projects are planned and developed, there is a need to perform a value analysis of geographically diverse sites on the efficiency of the overall wind plant.

Milligan and Artig [23] use hourly wind-speed data from six geographically diverse sites collected by the Minnesota Department of Public Service to provide some insight into the potential benefits of disperse wind plant development. We provide hourly wind power from each site to an electric reliability simulation model. This model uses generating plant characteristics of the generators within the state of Minnesota to calculate various reliability indices. Because data on wholesale power transactions were not available, these transactions are not included in the analysis. Although this work does not directly address the capacity credit issue, it is based on maximizing reliability, which is the basis of the commonly used ELCC. Therefore, the methods described here can be easily adapted to analyze the capacity credit of geographically disperse wind power plants.

The Potential Benefit of Wind Power Plant Geographic Diversity

Wind can be described as a stochastic process. As such, the power output from a wind power plant can vary substantially over time, and it is not controllable in the same way as conventional power plants. During lulls in the wind, other resources must supply electricity. If geographically diverse wind sites can be chosen in such a way as to minimize the number or extent of wind power lulls, this can be beneficial, because conventional resource use can be reduced accordingly. The extent of this reduction will influence the magnitude of fuel cost, O&M, and other costs to the utility. Of course the increase in wind power generation also creates other costs, such as O&M.

One of the first comprehensive studies to address the geographical diversity of wind plants was done by Kahn [24], who used California wind and utility data. He found that reliability does increase as a function of geographic dispersal, but the geographic wind diversity and the barrier of large wind plant penetrations relative to the conventional generator mix limit this increase. Kahn also points out that wind sites that are un-correlated will generally provide better combined reliability than sites that are highly correlated, in absolute value. However, Kahn's analysis ignores the fact that two or more wind regimes with significantly different time-scale properties can both provide the same correlation with utility load. A study by Brower [25] found some benefits to distributed wind development in Minnesota, but the benefits were somewhat constrained by the relatively high wind-speed correlation between wind sites.

State of Minnesota Data Collection Project

The wind resource data used in this study were collected through the Minnesota Department of Public Service's (DPS) wind resource assessment programs and the DPS/U.S. Department of Energy (DOE) Tall Tower Wind Shear Study. The DPS has conducted wind resource assessment since the early 1980s, providing utilities, developers, and other interested persons with wind data collected at sites around the

state. Since the programs began, DPS has expanded and improved the data-collection process by adding new monitoring sites and more sophisticated equipment.

The monitoring sites that provided the data for this paper are equipped with cellular data loggers that automatically send the collected information to a base station computer located in DPS offices. These sites use existing communication towers and have monitoring levels at 30, 50, and 70 meters (m) above ground level. Two anemometers are mounted at each level, one on each side of the tower. This configuration has several advantages. It reduces the wind shadow effect the tower would have on the data if only one anemometer were used at each level; it provides a degree of redundancy at each level so the failure of one sensor does not eliminate the data collection at that level; and it provides the opportunity to do sensor-to-sensor calibration and helps diagnose potential sensor problems. Each tower is also equipped with wind vanes at the 30-m and 70-m levels. In addition to the internal logger temperature, some of the sites are equipped with external temperature probes mounted approximately 4 m above ground level.

In 1996, the DPS, in co-operation with DOE, installed four advanced monitoring sites as part of the DPS/DOE Tall Tower Wind Shear Study. These sites use existing communication towers and have monitoring levels at 10, 30, 40, 50, 60, and 70 m above ground. As with the sites described above, each monitoring level has two anemometers, one mounted on each side of the tower. Each tower is also equipped with wind directional sensors at the 10-, 30-, 60-, and 70-m levels.

Wind Site Selection

Milligan and Artig chose six Minnesota wind sites to represent the diversity of climatology in the state. The sites selected are Alberta, Becker City, Brewster, Crookston, Currie, and Luverne. They appear in Figure 30, and are identified by the first two letters of the respective site name. One unique feature of the Minnesota DPS data collection effort is that wind data is collected to a height of 70 m. This made it possible to use the power curve of a modern utility-scale turbine at a hub-height of 65 m. A comparison of output at different hub-heights appears in Table 3. The research used wind data for one year beginning in November 1995.

	Correlation to Load	50m	60m	65m
Alberta	-0.0135	.287	.320	.332
Becker	-0.0436	.287	.311	.322
Brewster	-0.0395	.333	.362	.374
Crookston	-0.0035	.270	.297	.309
Currie	-0.0539	.362	.388	.402
Luverne	-0.0317	.320	.343	.358



Figure 30. Minnesota wind sites

For many utilities, simply maximizing wind energy capture during the system peak will not necessarily result in maximum benefit because the generating units' ramp-rates and minimum-run levels may not allow the wind power to be fully utilized if there is significant variation in the timing of wind power delivery to the grid. In addition, it is possible that the ability to accurately forecast hourly wind speeds and the consequent hourly wind power output is somewhat impaired by wind sites with high hourly variability. For these reasons, it may be to the utility's advantage to install wind plants and select among sites in such a way that hour-to-hour output variations between wind plants are reduced, while still obtaining as much wind power output as possible during the peak period. A further benefit, not studied here, may be related to the ability to improve forecast errors across multiple wind sites compared to forecast errors at a single site.

Figures 31 and 32 illustrate some potential benefit from different sites. The graphs are based on real wind data that was used to calculate the output of fictitious wind power plants, each with 100 MW of installed capacity. The hourly wind power is calculated by taking actual wind-speed data and calculating power output based on a utility-scale wind turbine, as described below. Figure 31 shows a 48-hour period for three of the sites used in the study: Alberta, Currie, and Luverne. During the first day, the highest output comes from the Currie site, and Alberta and Luverne show low power output. The second day, wind power output at Currie is low, whereas power output at both Alberta and Luverne reach maximum rated output (less losses). Power output at Luverne drops about 4–5 hours earlier than at Alberta late in the second day. Figure 32 shows another view of multiple-site wind power output. This 24-hour period shows a general correlation between sites, yet one can also detect time lags of 2–3 hours early and late in the day. These variations are somewhat typical of the data used in this study and illustrate the potential benefit of geographically diverse wind power plants.

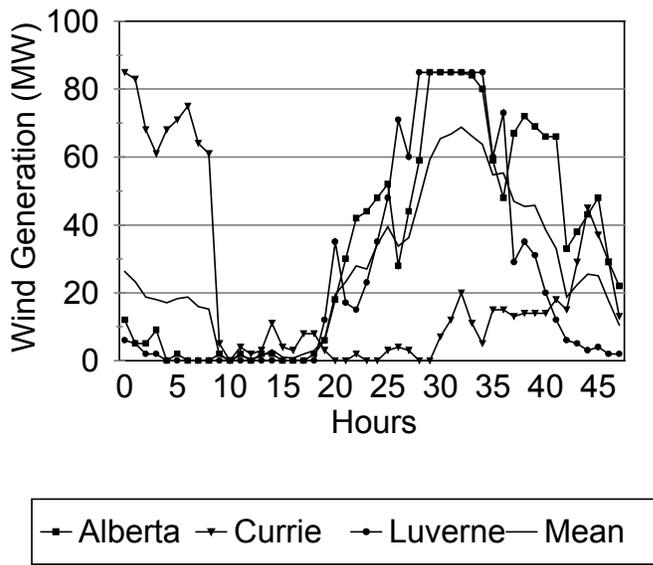


Figure 31. Wind power output over a 2-day period from multiple sites

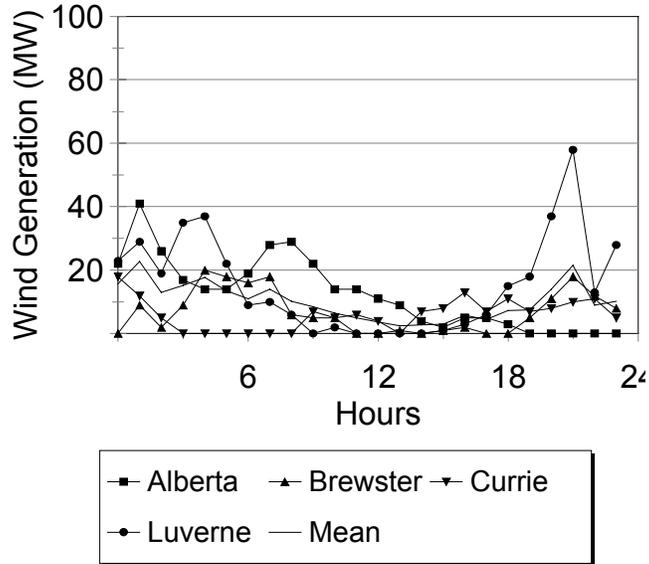


Figure 32. Wind power output over a 1-day period from multiple sites

Variations within a given wind plant are well known and provide some measure of diversity even within a relatively small geographic area. The extent to which this intra-site variability would influence geographic optimization is unclear. However, some of this intra-site variability is caused by local turbulence in what is known as the turbulent scale of the relative spectral intensity of the wind-speed, and is described by Stull [26]. For electrical reliability, we must be especially concerned with the synoptic scale variation, generally less than one cycle per hour. This allows us to choose among several potential wind-plant sites so that generator reliability is optimized over the year, or other relevant time period.

Modeling

Electrical reliability is a function of customer demand and the various generator characteristics. Utilities experience a pronounced peak period, often several hours of the day during a particular season. A utility can sometimes dramatically increase its generator reliability by installing peaking units to generate power when it is most needed: during the peak hours. Even though these peaking units might be available at night, their availability at night would likely have a negligible effect on system reliability. Likewise, as we have already seen, a wind plant that delivers a significantly higher annual energy output does not necessarily contribute significantly to system reliability. What is needed is for the wind output to occur at times of otherwise high-risk periods during system peak.

Kahn's discussion centers on the statistical correlation between the various wind-plant sites. Sites with high positive correlation will provide higher output during the same time periods, whereas sites with high negative correlation will be complementary. When the first site is providing a high level of electrical output to the grid, the second site will likely be idle. Conversely, when the first site is not producing electricity, it is likely that the second site is.

Figure 33 shows a normalized load duration curve for the state of Minnesota, which is part of the input to the Elfin model.

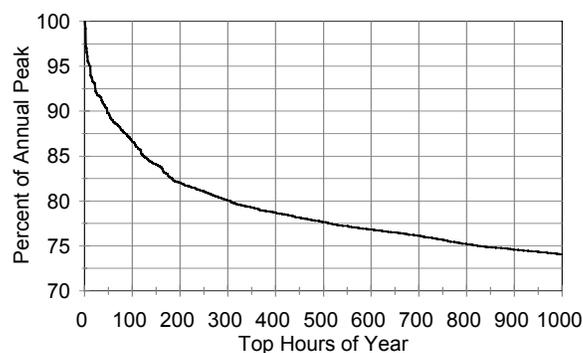


Figure 33. Minnesota annual load-duration curve

The precise form of the optimal wind-plant reliability problem may vary according to the customer loads, wind sites, and characteristics of the utility system. For example, a combination of wind regimes that all exhibit diurnal variation may provide the utility with an opportunity to select a combination of sites that together have a high probability of offsetting system peak requirements. Such a scenario might involve calculating various probability levels of wind generation during several peak hours every day during the peak month(s). Alternatively, when the wind resources do not follow a pronounced diurnal pattern, the utility might be more interested in looking at the overall probability levels of wind generation during the month, without necessarily allowing for a repetitive daily pattern in wind generation.

Choosing the best combination of wind sites can be done with a number of objective functions, depending on what the decision-makers believe is most relevant. Among these are (1) least-cost combinations of wind sites, (2) wind sites that minimize load swings during system peak, or (3) most reliable wind sites. The goals of least-cost production and most-reliable production are usually not consistent with each other. Reliability must be traded off against cost because a perfectly reliable system (if one were to exist) would not be cost-effective. Likewise, a least-cost solution might result in a generating system that does not possess sufficient reliability. However, if plausible estimates of outage costs are available, one could perform an optimization using the standard generation expansion options available in the model. For this analysis, Milligan and Artig pursued an optimization based on reliability. Their method can be easily extended to capacity credit or other parameters of interest. (Milligan and Factor [27] applied an extension of this procedure utilizing both a reliability and economic benefit optimization).

Marginal Reliability Method

The first method is based on traditional marginal analysis. Although this approach is widely used in economics, it is difficult to apply to the problem at hand. The difficulty is that the marginal reliability of a wind plant is potentially different in each hour of the year, depending on customer load and other generating resource availability. For example, we could calculate the reliability of wind power plants at Brewster and Currie. If Brewster appears to be the best choice based on its marginal reliability, we can add 100 MW. However, now that we have 100 MW at Brewster, the marginal reliability of Currie will be very different than during the original comparison with Brewster. Milligan and Artig used this method, and compare the results to those of the preferred methods, below.

Without specific price information on wind development and production at these sites, Milligan and Artig assume that the installed cost in \$/kW or \$/MW is the same at all sites. However, there is a difference in efficiency and reliability between sites because of the different winds experienced at each of the six different locations. The reliability level can be described as a function of installed MW at each site:

$$c_i = \Phi_i(x_i) \tag{9}$$

where c_i = reliable capacity, as measured as a function of $1/r$, r is the reliability measure of choice: either LOLE or ENS, x_i is the rated installed power capacity of the wind plant, i is the subscript of the wind plant location, and $1 \leq i \leq 6$. Finding the optimal mix of resources, assuming the same price/kW at each site, implies that the quantity of wind resources is deployed up to the point at which the marginal products of each site are equivalent (see Varian [28]). Written in terms of partial derivatives we have

$$\partial \Phi_i(x_i)/\partial x_i = \partial \Phi_j(x_j)/\partial x_j \tag{10}$$

$\forall i$ and j such that $1 \leq i, j \leq 6$.

Figures 34 and 35 show the reliability curves for two sites, selected to illustrate variations in reliability. The y-axis shows reliability, as measured by the marginal energy reliability index, which is the marginal ENS, scaled to the interval (0,1) for convenience. From the diagrams, it is clear that Brewster is the more desirable site based on its marginal reliability. Becker was not chosen by this (or any other) algorithm.

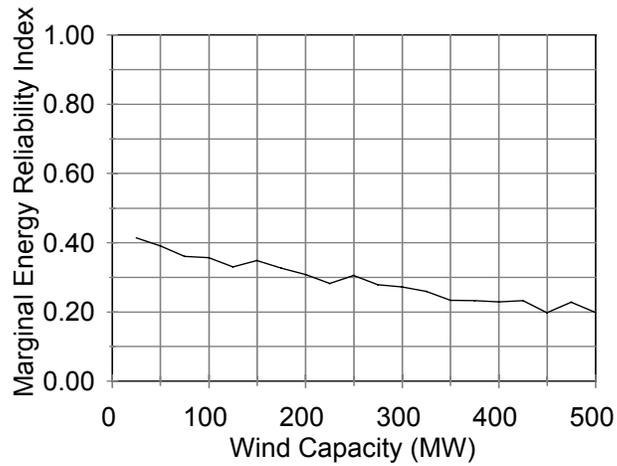


Figure 34. Marginal reliability curve for Becker

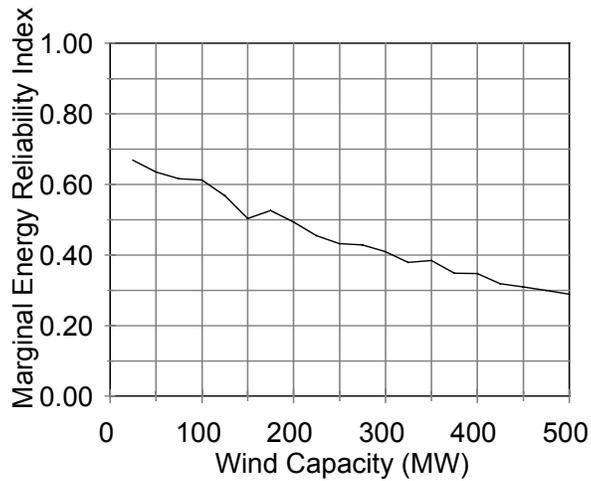


Figure 35. Marginal reliability curve for Brewster

For each of the six wind sites, a series of model runs were executed with Elfin, adding 25 MW at a time from each site until the maximum of 500 MW was “built” at each site. The authors then chose the combination of sites that satisfied equation (10), using the marginal ENS as the reliability measure, given the constraint of 500 MW of total installed wind capacity. The optimal choice from this method is 75 MW at Alberta, 150 MW at Brewster, 125 MW at Currie, and 150 MW at Luverne. The algorithm did not select either the Becker or Crookston sites because the associated reliability curves are much lower than those of the four selected sites.

Optimization with Elfin

This approach uses the production-simulation/reliability model in such a way as to do a step-wise modification of the net remaining load after each incremental wind plant is built. This method could perhaps be best understood by referring to Figure 36. There are two possible variations of this approach. The first uses LOLE as the optimization parameter; the second variation uses ENS. The step in the diagram that illustrates the choice of site with the best reliability parameter instructs us to choose the wind site with the *lowest* parameter, either LOLE or ENS, because higher values of LOLE and ENS represent less reliable systems. The step entitled “Build an X MW wind plant” implies that the incremental plant size to simulate building can be varied. In our case, X = 25 MW was chosen as a reasonable trade-off between accuracy and model run-time. Smaller values of X might be more accurate, although given the relative scale of X to the hourly loads, that is unlikely. Large values of X compromise the results because the optimization algorithm is restricted to large increments of wind capacity, possibly overshooting a better mix of sites.

Do until desired wind capacity is built

- *Calculate reliability parameter for X MW at each wind site*
- *Choose wind site with best reliability and “build” an X MW wind plant at this site*

Figure 36. Optimization algorithm

The optimization process proceeds as follows. This discussion focuses on using ENS as the reliability parameter, but the process is the same when LOLE is used as the optimization parameter. First, Elfin is run without any wind plants. The next step is to run Elfin for a block of 25 MW of installed wind capacity at each site, separately. The ENS calculation from each site is compared, and the site with the lowest ENS (best reliability) is selected. The process then simulates the building of 25 MW of wind capacity at the chosen site, and this becomes the new base case. The process is repeated, running Elfin for each site combined with the chosen site from the previous step. Twenty-five MW is chosen from the site with the best ENS, and the process is repeated until all 500 MW of wind has been installed. The algorithm then simply counts the number of 25-MW increments of wind plants added at each of the sites, and that is the result.

The LOLE optimization selects 250 MW at Brewster, 225 MW at Currie, and 25 MW at Luverne. The ENS optimization selects 450 MW at Brewster, 25 MW at Currie, and 25 MW at Luverne. The results from this set of optimizations appear in Table 4. However, these results do not tell the whole story. When we examined the selection part of the optimization, they found that there were often extremely small differences in either LOLE or ENS between the chosen site and the second or third runner-up. This issue is discussed further in the next section.

	Alberta	Becker	Brewster	Crookston	Currie	Luverne
MW Built by ENS Optimisation	0	0	450	0	25	25
MW Built by LOLE Optimisation	0	0	250	0	225	25

Inter-Annual Variations and Uncertainty

The wind-speed data used for this study were collected by anemometers mounted on a single tower at each of the six sites that were analyzed. Using a power-curve for a modern wind turbine, hypothetical power output was calculated, after accounting for wake effects and mechanical and electrical losses. If 25-MW clusters of wind turbines were built on any of these sites, however, each turbine would respond to somewhat different winds, depending on the terrain and micro-scale meteorological events. Therefore, we are forced to accept the proposition that each time-series of wind speeds represents one of many possible series. Although it is possible that each of these meteorological towers has been placed in a “representatives” location for the overall site, we have no assurances that this is indeed the case. This implies that the precise calculations from our models are based on somewhat imprecise data.

In previous work, we have also been somewhat skeptical of modeling that does not explicitly take inter-annual wind speed variations into account. Because of data constraints, Milligan and Artig were not able to perform a full analysis of the underlying time-series properties from multiple years of data at each wind site, although that would be our preferred approach (although in a later work [29] we were able to use several years of data). This would allow the use of sequential Monte Carlo runs with Elfin, resulting in probability distributions of the reliability measures for each wind plant and combination of plants. Such an analysis would allow for the explicit accounting of the underlying probability distributions, so as to help the decision-maker assess the impact of these variations. In the absence of a more detailed analysis, Milligan and Artig applied a technique borrowed from fuzzy logic (such as Monteiro and Miranda, [30] and Pereira [31]) to this problem. Using fuzzy logic allows the modeling to incorporate the uncertainty associated with the issues discussed above. Milligan and Artig hypothesized that the LOLE and ENS measures obtained from the Elfin optimization are fuzzy values, with variations ranging up to $\pm 0.5\%$ of the calculated value. The choice of this value range is based on the partial results of the optimization runs. As we examined the reliability values of the best site compared to the runners-up, there appeared to be a clustering of reliability values very close to the optimal values, whereas the least optimal plants’ reliability values were significantly worse. In our judgement, the choice of 0.5% was a reasonable one, based on the data. Without specific probability distributions, we hypothesize that the reliability measures are distributed uniformly on this interval, which is similar to other approaches using fuzzy analysis. The selection decision portion of the optimization algorithm was modified to select not only the best single site, but any site whose optimization parameter is within some small distance of the best choice. Since there is no *a priori* knowledge of which fuzzy value is best, the analysis used step-wise increments from 0.0% up to 0.5% of

the differences in the reliability measure and averaged the cases. This amounts to choosing a 25-MW block of installed wind capacity whenever

$$c_p(1 - \varepsilon) \leq c_i \leq c_p(1 + \varepsilon) \quad (11)$$

where c_p is the reliable capacity of the best site, ε is the fuzzy parameter expressed as a decimal, and c_i represents the capacity of plant i , $1 \leq i, p \leq 6$, and $p \neq i$. The results of this approach appear in Table 5.

	Alberta	Becker	Brewster	Crookston	Currie	Luverne
MW Built by Fuzzy ENS Optimisation	38	0	215	0	120	127
MW Built by Fuzzy LOLE Optimisation	65	0	155	0	173	108

The preferred method is the fuzzy ENS approach because ENS represents the area under the load probability distribution, whereas LOLE represents the height of the tail. Depending on relative costs of purchasing on-peak capacity and energy, a utility could use whichever method is most appropriate.

Comparison of Results

A summary of results appears in Figure 37. Each bar in the diagram represents a single method. The first three methods provide similar results: the best combination of sites excludes Becker and Crookston. In each of these cases, the lowest recommended capacity is at Alberta. Brewster and Currie are both recommended in the range of 125 MW to about 200 MW, and Luverne's share ranges from about 100 MW to about 150 MW. The right side of the graph shows how unstable the results can be when a deterministic approach is used. A small amount of capacity at Luverne is chosen in both cases, but there is clearly a very large difference in the capacity recommendations for Brewster and Currie. This disparity is caused by the extremely close reliability values that were often found among the runners-up. In this case, Brewster and Currie were very close in both the LOLE and ENS reliability measures, so small differences between these measures altered the relative ranking of the sites. This is one reason for applying a method that recognizes the role of uncertainty in the modeling. The method of choice is the fuzzy ENS approach because ENS typically provides a more robust measure of reliability, in general, than does LOLE, and is therefore more likely to be stable over short variations in load and generator parameters.

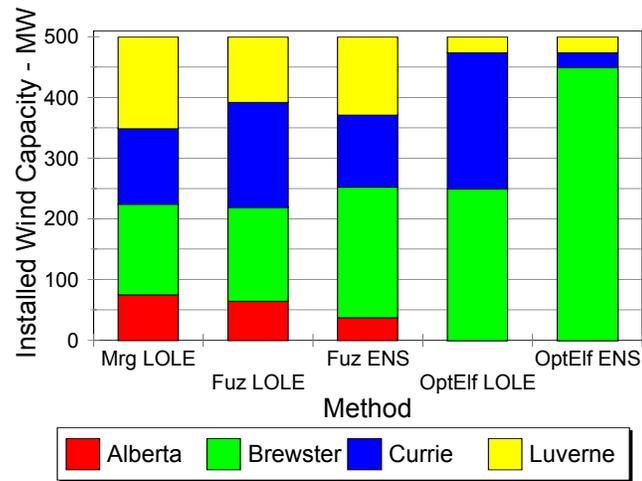


Figure 37. Comparison of methods

Summary

Electricity production simulation and reliability models can be applied to the problem of selecting among competing sites for wind generators. However, using these models must be tempered with some judgement. The wind sites analyzed by Milligan and Artig exhibit some overall correlation, but also provide some benefit to the overall system reliability because of time lags in hourly generation. We believe that the fuzzy ENS analysis provides the best way to analyze such problems.

Several additional factors could be introduced into future studies. First, given additional intra-site data, the results would be more accurate. Second, these results are sensitive to the specific load and generator characteristics used by the model. Additional data on wholesale power transactions from the state of Minnesota would improve the accuracy of these results. Constraints in the transmission system and power flow have not been considered here, but it would be important to analyze these factors before embarking on the installation of a large, geographically diverse wind power system.

Conclusions

The conclusions we can draw from this paper can be divided into two areas: modeling results and modeling methods.

Modeling Results

This paper describes modeling results for several utilities, wind sites, and wind turbines. In many cases, both an LDC model and a chronological model were used; in other cases, only an LDC model was applied. The first general observation is that, in all cases, modeled wind power plants have economic value—both energy benefits and capacity benefits were found to be significant. Wind power plants are factored into the dispatch mix by a simple, rational strategy: if there is wind capacity online, it is used. Although it may seem as though the wind power plant is a base-load unit, it will not displace base-load generation. Instead, high-dispatch-cost units that are on the margin are displaced by wind power.

The energy value of wind power plants is highly dependent on the utility, wind turbine performance characteristics, and wind site. Because the wind power displaces power generated by marginal units, the value of power displaced will vary throughout the day. During low-load periods, a marginal generator typically has a lower fuel cost than during the system peak. Therefore, the timing of the wind power has an important influence on the value of energy that is displaced. Wind sites that are highly correlated with load will have a higher energy displacement value because higher-cost energy is displaced during the peak period.

The capacity value can be calculated in the same way for a wind power plant as for a conventional generator. Although only two NERC regions explicitly deal with wind capacity accreditation, results from both methods show a capacity value that varies throughout the year. Wind forecasts can also have an important influence on operational capacity value of wind plants, because accurate forecasts help schedulers to commit slow-start units in an optimal fashion, making full use of the availability and timing of the wind power.

Using a reliability-based measure of capacity credit, the ELCC, it is apparent that wind plants have a capacity credit that varies with the characteristics of the wind site and utility characteristics. In all cases examined in this report, wind had significant capacity credit value, although this value is a percentage of the wind plant nameplate capacity. What is the implication about installing backup capacity for wind power plants? In our judgement, the answer is the same as for conventional power plants. A conventional power plant may be on forced outage during a system peak. Although this event normally has a low probability of occurrence, the unit might not be able to contribute capacity during system peak. However, the capacity rating for the plant will not normally be zero. We can use a measure such as ELCC as the capacity value of the plant, so as to incorporate the expected forced outage probability into the capacity rating. Should the plant be on outage during the peak, the shortfall would be made up from system spinning reserve, quick-start units, spot purchases, or by other available units that can be committed in time to cover the peak. Historically, reserves have been provided on a system-wide basis, not on the basis of an individual power plant. The relative burden that each power plant puts on the reserve requirement is a function of the plant capacity and its availability rate.

To maximize the reliability contribution of wind capacity to the generating system, our results suggest that development occur at several sites rather than at a single site. The analysis does not consider the effects on the transmission system or the possibility of reducing wind forecast errors of geographically disperse development. If a single large wind site were to be developed, it would likely require a dedicated transmission line that would have a load factor approximately equal to that of the wind power plant. A line outage would result in the loss of all wind generation. Conversely, with multiple wind sites, it may be possible to use existing transmission, attaining a higher line-load factor if the line can be used for other purposes. A line outage would not necessarily result in the loss of all wind capacity.

Modeling Issues

It is apparent that chronological models provide a more detailed, accurate look at generation costs than do LDC. Having said that, LDC models, as represented by Elfin, remain very powerful tools. Overall cost estimates between these two models were very close. It is likely that data quality and availability, along with other modeling assumptions, play a more important role in accurate modeling than does the style of the model. LDC models are not capable of performing daily analyses, which are important as we move from planning issues to operational issues.

Another issue that affects accuracy is whether the model uses a piecewise representation of the load curve or a cumulant method. This poses a trade-off between accuracy and execution speed. Accuracy losses are more apparent with systems that have a high reliability because it is more difficult to represent the tail of the probability distribution, as the probability of loss-of-load becomes very small.

Computing ELCC can be somewhat time consuming and requires a significant level of modeling and data requirements. In the absence of these, calculating wind plant capacity factor over the top 10% of load hours can often provide a reasonable first-cut estimate of capacity credit. Although this method may not always perform accurately, it is much easier to apply than a set of production-cost/reliability model runs.

Using a single year of wind data to estimate capacity or energy value is somewhat risky because of possible inter-annual variation in the wind resource. Using Sequential Monte Carlo, it is possible to calculate measures of dispersion, such as the standard deviation, of capacity and energy benefits. It also might be possible to correlate wind site data with data at a nearby site that has a longer-term record. Sequential Monte Carlo can also be used to provide probability distributions for any model output of interest, such as total cost, cost of specific generators, and so on. Simulated data is certainly not a substitute for real data. As more multi-year wind data sets become available, analyses can incorporate the data instead of simulating it.

A technique of sampling from multiple wind realizations can help reduce model run-time. Although the REPA process does a reasonable job of capturing the variation in wind capacity credit, other sampling methods might do a better job of representing the ELCC distribution and the full variation of wind energy output found in the actual data set. In my judgement, reduction techniques hold great promise for providing significant information while limiting model runtime requirements. For example, it may be possible to model ELCC as a function of various parameters using some form of regression or time-series modeling. The REPA sampling would then have more guidance as to which series would best represent the ELCC variation. Of course, this could be applied to other model outputs as well.

When multiple wind sites are available for development, the optimal mix of these sites can be calculated according to a number of objective functions. If capacity and energy outage costs are known or can be estimated, generation-expansion model runs that use a dynamic programming procedure might be useful. In the absence of these costs, an iterative method based on fuzzy ENS can be used to calculate which sites should receive development. The procedure outlined here has considerable flexibility, so it can be adapted to specific applications as appropriate. This flexibility makes it possible to analyze competing objectives, such as minimum generator cost vs. maximum reliability. Decision-makers can then choose to combine the results in an attempt to meet both objectives.

Because wind is an intermittent resource, modeling approaches based on probabilistic and fuzzy set theory are powerful ways to analyze possible outcome ranges. These approaches have been applied to various wind analysis projects in Europe and North America. Power systems analysts are becoming more adept at quantifying risks—risk of generator failure, transmission failure, fuel cost variations, and others. Another form of risk that is introduced by wind power plants is the risk of no wind power output. As we have seen in this paper, there are ways of incorporating these risks into the analysis. As the use of wind power plants increases throughout the world, techniques for addressing these issues will undoubtedly increase in sophistication. Although the restructuring landscape is still undergoing change, increased competition among electricity suppliers will increase the demand for modeling tools and approaches that can help assess risks, reliability, and intermittence.

Assessing the Modeling Frameworks

It is not the intent of this report to weigh the costs and benefits of the two models, but rather to comment on the state of wind-power plant modeling as represented by the two models used for this research. Both P+ and Elfin allow the analyst to perform detailed analysis of wind power plants in an excellent modeling framework. Although the chronological framework can answer more detailed questions about optimal generation with wind power plants, both models suffer from the same weakness—the ability to simultaneously consider hourly variations in wind power with a probability other than certainty. Simple approaches, such as using the standard deviation of wind power during the period with an appropriate probability distribution, would help. On the other hand, sequential Monte Carlo, although computationally expensive, can provide a detailed look at variations, and can be weighted according to probability distributions derived from the simulated series, as performed by the REPA analysis above. Although REPA certainly has room for improvement, the approach is sound, and provides a view of probabilistically weighted outcomes and various representative cases. One of the appeals of Sequential Monte Carlo and REPA is the flexibility to use the method with any model.

Future Research

As we move towards a time that will likely feature a restructured electricity market and vastly improved wind generating technology compared to the past, fundamental questions about the operational impacts of wind power plants on the power system need to be fully answered in the public domain. Does a large wind plant increase or decrease the high-resolution variation in net load (after accounting for wind power)? What are the effects on unit ramping? If these effects are significant, what mitigation techniques can be applied? How much short-term smoothing occurs as the number of wind turbines within the wind plant increases? Can control strategies for single variable-speed machines be adapted for larger clusters of turbines, minimizing losses and short-term fluctuations? How is the system reliability affected by the penetration rate of geographically disperse wind plants? And how much more penetration can we get with additional disperse sites?

References

- [1] Milligan, M. Modeling Utility-Scale Wind Power Plants. Part 1: Economics. *Wind Energy* 1999, **2**, 167–193.
- [2] Wan, Y. An Analysis of Current Utility Practice of Rating and Reporting Wind Generating Capacity. Working paper. Golden, Colorado: National Renewable Energy Laboratory, 1997.
- [3] Milligan, M., A. Miller, and F. Chapman. Estimating the Economic Value of Wind Forecasting to Utilities. *Windpower '95 Proceedings; March 27–31, 1995*. Washington, DC: American Wind Energy Association, 1995.
- [4] Landberg, Lars. Operational Results from a Physical Power Prediction Model. *European Wind Energy Conference Proceedings*. March 1–5, 2000; Nice, France.
- [5] Milligan, M. A Chronological Reliability Model to Assess Spinning Reserve Allocation to Wind Power Plants. Accepted for the *European Wind Energy Conference*, July 2-6, 2001. Copenhagen, Denmark. To appear.
- [6] Panel session: “Maintaining Reliability in a Competitive Environment.” *Probabilistic Methods Applied to Power Systems 5th International Conference*, Vancouver, British Columbia, Canada, September 1997.
- [7] Milligan, M. “Measuring Wind Plant Capacity Value.” *WindStats Newsletter; Vol. 9, No. 1*; Knebel, Denmark. NREL/TP-441-20493. Golden, Colorado. National Renewable Energy Laboratory, 1996.
- [8] Milborrow, D. “Struggling With Capacity Credit Definitions.” *Windstats Newsletter; Vol. 9, No. 2*; Knebel, Denmark, 1996.
- [9] Milligan, M. and B. Parsons. “A Comparison and Case Study of Capacity Credit Algorithms for Intermittent Generators.” Presented at Solar '97, Washington DC, April 27-30, 1997. NREL/CP-440-22591. Golden, Colorado: National Renewable Energy Laboratory, 1997.
- [10] Kahn, E. *Electric Utility Planning & Regulation*. American Council for an Energy-Efficient Economy. Washington, D.C., 1991.
- [11] Billinton, R. and R. Allan. *Reliability Evaluation of Power Systems*. New York: Plenum Press, 1984.
- [12] Milligan, M. “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, 1996.
- [13] Percival, D. and J. Harper. *Value Analysis of Wind Energy Systems to Electric Utilities*. SERI/TP-732-1064. Golden, Colorado: Solar Energy Research Institute, 1981.
- [14] Bernow, S., B. Biewald, J. Hall, and D. Singh. Modeling Renewable Electric Resources: A Case Study of Wind. Tellus No. 91-187. Tellus Institute: Boston, Massachusetts, 1994.

- [15] Billinton, R., H. Chen, and R. Ghajar. "A Sequential Simulation Technique for Adequacy Evaluation of Generating Systems Including Wind Energy." IEEE/PES Winter Meeting, January 21–25, 1996. Baltimore, Maryland, 1996.
- [16] Milligan, M. and M. Graham. *An Enumerated Probabilistic Simulation Technique and Case Study: Integrating Wind Power Into Utility Production Cost Models*. NREL/TP-440-21530. Golden, Colorado: National Renewable Energy Laboratory, 1996.
- [17] Marnay, C., and T. Strauss. *Variance Reduction in Monte Carlo Chronological Production-Cost Modeling*. Sacramento, California: California Public Utilities Commission, 1990.
- [18] Milligan, M. "Wind Plant Capacity Credit Variations: A Comparison of Results Using Multiyear Actual and Simulated Wind-Speed Data." *Windpower '97 Proceedings; June 15–18, 1997*; Austin, Texas. NREL/CP-440-23096. Golden, Colorado: National Renewable Energy Laboratory, 1997.
- [19] Milligan, M. and M. Graham. An Enumerative Technique for Modeling Wind Power Variation in Production Costing. *Proceedings of the Probabilistic Methods Applied to Power Systems 5th International Conference; September 21-25, 1997*; Vancouver, BC, Canada. Vancouver, BC, Canada: BC Hydro. NREL/CP-440-22868. Golden, Colorado: National Renewable Energy Laboratory, 1997.
- [20] Milligan, M. and J. Kline. An Evaluation of Hourly Average Wind Speed Estimation Techniques. *WindPower '98*, Bakersfield, California. NREL/CP-500-24664. Golden, Colorado: National Renewable Energy Laboratory, 1998.
- [21] Hillier, F. and G. Lieberman. *Operations Research*. Holden-Day: San Francisco, 1974.
- [22] Box, G.E.P, and G.M. Jenkins. *Time Series Analysis: Forecasting and Control*. Holden-Day: San Francisco, 1976.
- [23] Milligan, M. and R. Artig. Optimal Site-Selection and Sizing of Distributed Utility-Scale Wind Power Plants. *Proceedings of the 21st Annual International Conference of the International Association of Energy Economics*; Quebec City, Quebec, Canada. NREL/TP-500-24312. Golden, Colorado: National Renewable Energy Laboratory, 1998.
- [24] Kahn, E. The Reliability of Distributed Wind Generators. *Electric Power Systems Research*. **2**. Elsevier Sequoia. Lausanne, 1979.
- [25] Brower, M. *Powering the Midwest: Renewable Electricity for the Economy and the Environment*. Union of Concerned Scientists. Cambridge, Massachusetts, 1993.
- [26] Stull, R. *An Introduction to Boundary Layer Meteorology*. Kluwer Academic Publishers. Norwell, Massachusetts, 1988.
- [27] Milligan, M. and T. Factor. Optimizing the Geographic Distribution of Wind Plants in Iowa for Maximum Economic Benefit and Reliability. *Wind Engineering*. **24**, No. 4, 2000. 271-290.
- [28] Varian, H. *Microeconomic Analysis*. W. W. Norton, New York, 1978.
- [29] Milligan, M. and R. Artig. Choosing Wind Power Plant Locations and Sizes Based on Electric Reliability Measures Using Multiple-Year Wind Speed Measurements. *Proceedings of the North American Conference of the United States Association for Energy Economics*. Orlando, Florida, U.S.A. NREP/TP-500-26724. Golden, Colorado: National Renewable Energy Laboratory, 1999.

- [30] Monteiro, C. and V. Miranda. Probabilities and Fuzzy Sets in the Market Evaluation of Renewable Energies: The Solar-GIS Experience. *Proceedings of the Probabilistic Methods Applied to Power Systems 5th International Conference; September 21–25, 1997*; Vancouver, BC, Canada. Vancouver, BC, Canada: BC Hydro, 1997
- [31] Pereira, J., J. Saraiva, and V. Miranda. Combining Fuzzy and Probabilistic Data in Power System State Estimation. *Proceedings of the Probabilistic Methods Applied to Power Systems 5th International Conference; September 21–25, 1997*. Vancouver, BC, Canada. Vancouver, BC, Canada: BC Hydro, 1997.

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.			
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE March 2002	3. REPORT TYPE AND DATES COVERED Technical report	
4. TITLE AND SUBTITLE Modeling Utility-Scale Wind Power Plants Part 2: Capacity Credit		5. FUNDING NUMBERS WER1.3010	
6. AUTHOR(S) Michael R. Milligan			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) National Renewable Energy Laboratory 1617 Cole Blvd. Golden, CO 80401-3393		8. PERFORMING ORGANIZATION REPORT NUMBER NREL/TP-500-29701	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES			
12a. DISTRIBUTION/AVAILABILITY STATEMENT National Technical Information Service U.S. Department of Commerce 5285 Port Royal Road Springfield, VA 22161		12b. DISTRIBUTION CODE	
13. ABSTRACT (<i>Maximum 200 words</i>) As the worldwide use of wind turbine generators in utility-scale applications continues to increase, it will become increasingly important to assess the economic and reliability impact of these intermittent resources. Although the utility industry appears to be moving towards a restructured environment, basic economic and reliability issues will continue to be relevant to companies involved with electricity generation. This paper is the second in a two-part series that addresses modeling approaches and results that were obtained in several case studies and research projects at the National Renewable Energy Laboratory. This second paper focuses on wind plant capacity credit as measured with power system reliability indices. Reliability-based methods of measuring capacity credit are compared with wind plant capacity factor. The relationship between capacity-credit and accurate wind forecasting is also explored.			
14. SUBJECT TERMS wind power plants; renewable energy; power system reliability; wind capacity credit; wind forecasting value; utility production costs; utility production-cost models		15. NUMBER OF PAGES	16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL