NREL/TP-440-21411 • UC Category: 1210 • DE96013072

Alternative Methods of Modeling Wind Generation Using Production Cost Models

Michael R. Milligan
National Renewable Energy Laboratory

Chok K. Pang, Granger Tam
P Plus Corporation

Prepared for
Conference of the Electricity Supply Industry
Kuala Lumpur, Malaysia

October 21-25, 1996

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

Prepared under Task No. WE617010

August 1996

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

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

Printed on paper containing at least 50% wastepaper, including 10% postconsumer waste

ALTERNATIVE METHODS OF MODELING WIND GENERATION USING PRODUCTION COSTING MODELS

Michael Milligan
National Renewable Energy Laboratory
1617 Cole Boulevard
Golden, Colorado 80401
USA

Chok K. Pang, Granger Tam
P Plus Corporation
20370 Town Center Lane
Cupertino, California 95014
USA

ABSTRACT

This paper examines the methods of incorporating wind generation in two production costing models: one is a load duration curve (LDC) based model and the other is a chronological-based model. These two models were used to evaluate the impacts of wind generation on two utility systems using actual collected wind data at two locations with high potential for wind generation. The results are sensitive to the selected wind data and the level of benefits of wind generation is sensitive to the load forecast. The total production cost over a year obtained by the chronological approach does not differ significantly from that of the LDC approach, though the chronological commitment of units is more realistic and more accurate. Chronological models provide the capability of answering important questions about wind resources which are difficult or impossible to address with LDC models.

1.0 INTRODUCTION

Utilities have been using production cost and generation expansion models for many years to evaluate the economics of future power production options. Uses for these models range from short-term budgeting to long-range planning. Early models were built on the important work of Baleriaux [1] and Booth [2], who developed an algorithm that combined hourly loads into a load duration curve (LDC). Under certain assumptions, the load duration curve may be interpreted as a probability distribution. The appeal of these LDC models was that they are computationally much less demanding than chronological models.

Chronological production cost models maintain the time chronology of the study period so that time dependent constraints such as unit minimum up and down times, ramp rates and unit commitment can be properly modelled. As computing power increased and became less expensive, chronological models received more attention. Thus more detailed production simulations can be performed, giving more detailed and accurate results [3].

The choice of model type (chronological or LDC) is particularly relevant for evaluating intermittent renewable resources such as wind power. Many potential effects of specific wind regimes on specific utility systems remain unknown, and production cost models can help answer some of these questions. For example, if wind power is introduced to a utility system, which generating units will be affected and by how much? What is the value of the displaced energy/capacity from the conventional units? Will the temporal variation in the wind power be comparable to that in the utility's native load? To properly answer these questions, we must be knowledgeable about the tools that do the measuring and how they treat intermittent resource modeling. The work outlined herein is motivated by the following question: Does the LDC modeling paradigm impose any significant shortcomings in measuring the value of intermittent resources such as wind?

This paper presents case studies of two utilities, each modelled with both an LDC-based model and a chronological model, to illustrate the differences and similarities that result when significant amounts of wind power are modelled on the system. Two wind regimes were used. Our intent is to comment on the differences between modeling paradigms rather than to judge the specific models.

The paper begins with a discussion of some general modeling issues that apply to intermittent resources such as wind, followed by a more detailed discussion of the model techniques used herein. We conclude with some selected case study results.

2.0 MODELS AND TECHNIQUES

The models used in the study are Elfin (developed by the Environmental Defense Fund) and P+ (developed by P Plus Corporation). Elfin is a LDC model while P+ is a chronological model. The major difference between these models lies in the algorithm used; there are several other differences that would result in different results even if the algorithms were to match. We attempted to minimize these differences when running the models.

2.1 Elfin Model

The Elfin model uses a representative week for each month, although the model can be run in 52-week mode. Elfin divides the chronological loads into different time subperiods specified by the user, such as peak and off-peak subperiods. For each of these subperiods, the model first constructs the piece-wise linear LDC, then performs the unit commitment and economic dispatch. Within each of these subperiods, the model assumes that each supply-side resource provides a constant level of capacity

Elfin allows for a large number of generator types, including both supply-side and demand-side resources, both firm and non-firm. There are several types of load modifiers, which can be applied as either demand-side or supply-side resources depending on whether unit commitment and loss calculations are to include the resource. Historically, intermittent generators, such as wind, have been modeled as load modifiers, because they are applied to the chronological loads prior to constructing the LDC.

Supply-side resources can be designated with either fixed or time-varying generation schedules. If a resource is not a load modifier, its capability is averaged over the time period, accounting for forced and planned (maintenance) outages, unless it is modeled with a series of forced outages rates. Generators can be subjected to certain constraints, most notably energy constraints for months, seasons, or years. A generator can be designated as "must-run," so that it runs irrespective of its cost, but otherwise generators are deployed in merit order.

2.2 The P+ Model

The P+ model actually contains several, separate program modules. The modules, PMONTH and PWEEK, are the chronological production costing modules and produce results for the hour, day, month, and year. These two modules use the same algorithm but have different time scales. PWEEK was the module used in this study.

P+ allows for various generation types, including thermal units, combustion turbines, hydro, energy storage, and transactions/interchanges. Commitment and dispatch occurs in sequence by type, beginning with transactions, observing fixed energy constraints, and dispatching in merit order. Committed units are dispatched first, and their minimum up-time and down-time constraints are enforced. It has three simulation options. The probabilistic option convolves each unit with respect to its forced outage rate after the initial dispatch [4]. The Monte Carlo option makes random draws according to the outage rate of each unit for its availability. This approach is computationally much more intensive, but is useful in estimating the mean and variance of the results. The third option is called "branching" and allows the user to specify a generator as the branching node. Each outage state of the branching node is simulated, followed by convolution of the other units as in the probabilistic option. This method allows the user to focus on outage effects of particular units while retaining the probabilistic approach for the other units.

Transactions can be specified on the basis of user-defined subperiods. The P+ model provides a variety of options for shaping a transaction with respect to peaks and valleys. Load modifications for demand-side management (DSM) or intermittent resources can be modelled as transactions. Thermal units are modelled using capacity, heat rate, fuel price, and availability data.

2.3 Methods for Modeling Wind Power

On an *a priori* basis, it is reasonable to assert the superiority of the chronological approach over the LDC approach for modeling wind power. However, it is not intuitively clear as to what the extent of the differences might be between these methods. Furthermore, many other factors will also influence the results. For example, in earlier work [5], we demonstrated a range of results from the same model for different methods of wind power modeling and wind data sampling. If such differences could be obtained using a single model, what are the implications of using more than one model? A larger question is, which model is closest to reality? Due to the partial predictability of wind, model "backcasting" must take the predicted level of wind power into account when looking backward in order to recognize the accuracy of hindsight over that of foresight [6]. There is no widespread agreement on model accuracy because many production cost and generation expansion models use different methods.

Given the range of results described in an earlier paper [5], one goal here is to begin looking at whether generalizations can be made with respect to the various LDC model results compared to chronological model results. However, there are several situations in which chronology may be important, the most significant of which may be the differences in unit commitment. In LDC models, committed units are fixed for the whole subperiod and they cannot be decommitted until the following period, which can be either a week or a month. Chronological models can commit or decommit a unit in any hour depending on need, and subject to its minimum up-time and down-time constraints.

Extreme short-term variations in wind power output could possibly cause excess cycling of other units which are required to support periods of low wind power output. During periods of high wind power output, these units would not be needed. If pushed to the limit, the unit cycling could conceivably become constrained by unit ramping constraints. If this were to occur, excess capacity would probably be required during periods of low-wind output because unit ramping constraints would not allow the reduction of conventional output quickly enough to compensate for the increase in wind power output. These periods of excess capacity would result in a reduction in the potential benefit of wind power.

Analysis of actual wind data indicates that large swings in wind power are possible, but they do not occur with high frequency. Additional analysis and data are required to draw more general conclusions. While these preliminary results justify a utility's interest in precise estimates of the effect of ramping on its load-following units, they suggest that the effects may not be severe but warrant further examination.

Most production cost models offer two broad methods of modeling intermittent resources such as wind. The first method is to treat the wind power as a load modification. The hourly wind power production is subtracted from the system load before performing commitment and economic dispatch. The advantage of this approach with LDC-based models is that it retains all of the chronological detail on the wind power. A disadvantage of the load-modification method is that it assumes the wind power output will be achieved with certainty, and there is no explicit accounting for the probability of higher or lower output.

The second method is modeling the intermittent resource as a supply-side unit. The wind is dispatched in merit order and convolved as with the conventional units. Wind power must be priced such that it would be dispatched at proper merit order. There is the intellectual appeal that wind power is treated in much the same way as the conventional supply-side units in the commitment/dispatch process, rather than as a load modifier. As such, forced outage rates can be used to incorporate the uncertainty of wind power.

The supply-side approach has several variations. The first is to provide the model with a value for expected wind power for the month, for example, along with a forced outage rate that is based on long-term wind availability during that period. The interpretation of this approach is that we are providing the production cost model with an expected level of wind power for the period, along with a measure of the variance. A second variation of the supply-side approach is to allow for time-varying forced outage rates. This method is called the "probabilistic supply-side" method and can be viewed as providing a time-dependent probability distribution that represents the availability of wind power [5]. It has the advantage of providing both a time-varying component and a probabilistic component to the wind power calculation. Greater accuracy can be achieved when a model allows for higher levels of temporal disaggregation and for more precise specification of the probability distributions. The relative performance of the two supply-side modeling variations, either with respect to each other or with respect to the load modifier approach, depends on the characteristics of the system that is modelled.

The two modeling methods, load modification and supply-side, are also available with chronological models. For example, P+ allows time-varying capacity to be modelled as a transaction, which results in the subtraction of wind power from the system load. After all transactions have been processed, the commitment and dispatch of the conventional units are determined. This approach is subject to the same criticism as with the LDC model – the conventional hydro-thermal system need only meet the modified load net of wind power. This tends to give more capacity credit to the wind power than actually exists.

Alternatively, the wind resource can be modeled as a supply-side unit with chronological models. We can specify both the time variation and probabilistic nature of the wind resource. Although P+ does not allow for variation in the probability values, it does allow the unit output level to change, given a fixed set of probabilities. Using this method, five points on a probability distribution curve are selected and the wind power level is allowed to change hourly, so the expected value of the wind power can be specified as a function of time. An advantage of the underlying chronological structure of the model is that ***hourly*** expected wind capacity is calculated and folded into the ***hourly*** calculation of the conventional units. At present, the Wind Power Simulator pre-processor (WIPS) model used in this work does not have the capability of calculating the inverse of the large number of probability distributions that would be required by the time-varying approach used by P+, so we were unable to test this approach.

2.4 Commitment and Economic Dispatch

The unit commitment methods used by these two models are different. The Elfin model requires the user to specify a commitment target for each subperiod, which can be a combination of a percentage of load and an adder to the load. The commitment order is based on the expected average variable cost of each unit at its most efficient operating level. The committed capacity is reduced by the forced outage rate, resulting in an expected capacity level for the unit. Elfin performs its unit commitment routine once for each period. The unit commitment is based on the peak load in the period and remains the same during the whole period. Quick-start units can be counted toward the commitment target.

Elfin's economic dispatch attempts to maintain consistency between chronologically adjacent subperiods by simultaneously dispatching across the subperiods. Chronology is only observed at the subperiod level, keeping the model from making illogical dispatch decisions such as dispatching a slow-start plant in the morning and afternoon but not during the midday. Dispatch proceeds in merit order, subject to must-run constraints. Undispatched portion of a committed unit is counted toward spinning reserve.

The P+ commitment logic for thermal and pumped storage units is conceptually simpler and logical, because the model proceeds chronologically. Fixed-energy transactions are dispatched first, followed by hydro and limited energy resources, and finally thermal and pumped storage systems. Slow-start units are committed and dispatched in accordance with their ramp rates. When committing a unit, it is subject to its minimum up-time and down-time constraints. The accuracy of the commitment process is quite good, because the target is a function of the chronological load. After commitment and dispatch have been calculated, the model checks to make sure that no contract fuel violations have occurred. Should a violation occur, P+ adjusts its commitment schedule within the fuel delivery and storage constraints specified by the user.

3.0 Case Study and Results

We set up a series of cases using two large utilities and two wind regimes. We calibrated the model inputs so that the results from the models matched as closely as possible, then ran various wind generation cases with each model. To hold the inputs constant between the models so that we could concentrate on algorithm differences, we used one representative week of load data per month in a year for both models. This would not be an approach taken by users of chronological models, but allowing the two models to use different system loads would introduce another source of variation. Certain power purchase and exchange contracts were simplified in both models to facilitate comparison between the various cases.

We provided both Elfin and P+ with 8,760 hours of wind generation for the year. Elfin performs a subperiod averaging routine so that each month's data are averaged into a typical week; it dispatches one typical week per month, and then scales the weekly values to the monthly values. P+ dispatches each typical week four times, using a full month of wind generation, which normally varies from week to week. Ignoring all other differences between the models, any variation in the results of the models is due to the constraints imposed by the LDC simplification.

Table 1: Resource Mix of Two Utility Systems

Resource Type	Utility U1	Utility U2
Nuclear	9.00%	9.08%
Oil/Gas Thermal	28.80%	31.35%
Hydro	15.60%	4.13%
Combustion Turbines	1.50%	2.20%
Purchases	40.00%	45.00%
Others	5.10%	8.25%

Table 1 shows the resource mix of the two utilities with a range of generator types and transactions. We explored several methods of modeling wind power in these two utility systems. For each method, we compared a base case with no wind power to cases with wind power simulated from two distinct wind regimes. The data for these hypothetical wind powers are based on actual wind data collected at two sites as part of the Department of Energy Candidate Site program and processed with the WIPS preprocessor model [7]. The first wind site is a high plains location (referred to as HP hereafter), and the second is located on a west coast mountain pass (WC). Power output was calculated on an hourly basis, accounting for electrical losses, wake effects, and mechanical availability. Curves showing the fit between the actual data distribution and two probability distributions, the Rayleigh and the Weibull, are shown in Figures 1 and 2. The wind power modelled for each system is approximately 4% of the system annual peak load (at rated capacity).

Two sets of load data were studied: a "high-load" case and a "low-load" case. The former provides us with a look at the two utilities when capacity is in short supply and peak commitment targets cannot always be met. The "low-load" case allows us to examine situations in which the utilities can meet the commitment target.

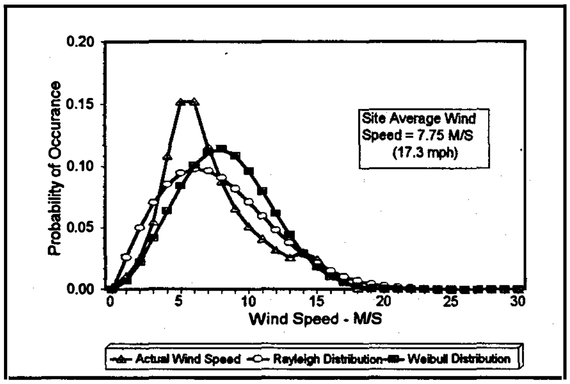

Figure 1: Wind speed distribution, HP wind site

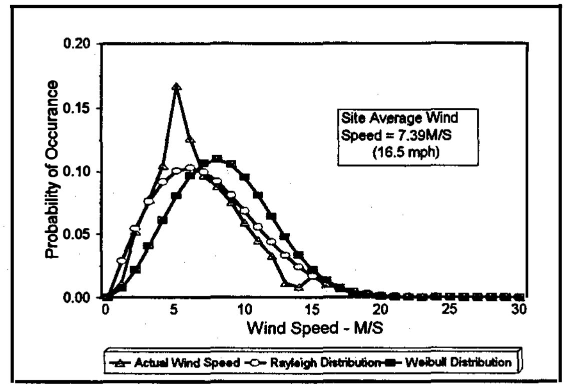

Figure 2: Wind speed distribution, WC wind

3.1 P+ Results

We begin with some of the P+ model results for the system peakday using the high-load scenario for Utility U1. Figure 3 shows the hourly generation by resource type for the no-wind case. For clarity, purchases have been taken out of the graph. The generation is plotted in (approximately) dispatch order, from the bottom up. On the bottom we see the base-load nuclear units, whose output is constant throughout the day. The "other" category includes a mix of other base-load units that also remain constant. The next level up are the gas/oil units, which make up a large contribution to the daily load swing, as do the hydro plants in the next level. At the very top we see the small contribution of the combustion turbines, which operate for approximately 6 hours around the peak.

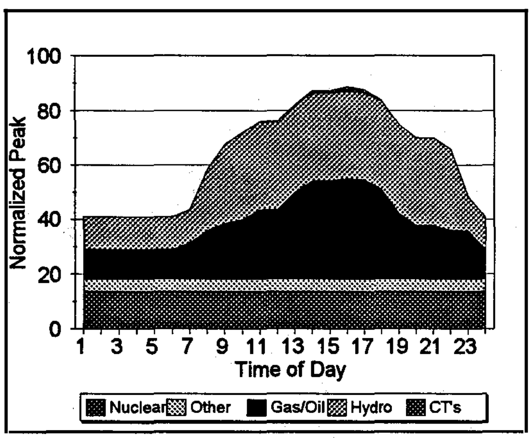

Figure 3: P+ peakday simulation for Utility U1, no wind

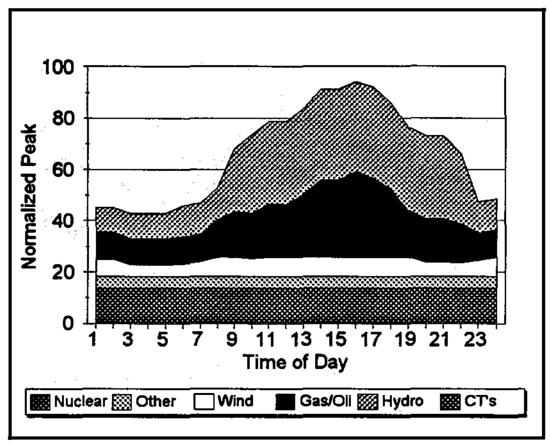

Figure 4: P+ peakday simulation for Utility U1, High Plains wind

Figure 4 shows the peakday effect of adding the HP wind power. The wind power is modelled as a load modifier. The wind power is graphed immediately above the base-load generation to help clarify the constant base-load output with and without wind. Figures 3 and 4 show that the units most affected by the introduction of wind are the intermediate and peaking units, and *not* the baseload units. The numerical simulation results show that the output of all baseload units is virtually unchanged by the introduction of wind power.

It is possible that, during extreme low-load periods and subject to operating constraints on the units, base-load output could decrease with the addition of wind power. The most expensive units are usually the first to be curtailed. In various cases, virtually *all* resources exhibited at least a small change when wind power was introduced into the resource mix. The only exceptions are those resources and/or transactions that are fixed. For a few transactions, the total energy did not change but the timing of the transaction sometimes shifted. The system marginal cost decreases with wind power. This caused a reduction in payments to the supplier whose outputs are priced at marginal cost, providing an additional cost benefit to the utility.

Figure 4 indicates that the wind resource is effective in reducing fossil fuel consumption. It illustrates a fortuitous combination of high system peak during a period of high wind power output. On the other hand, the WC wind regime generates a small amount of power during the early morning and during the daily peak [8].

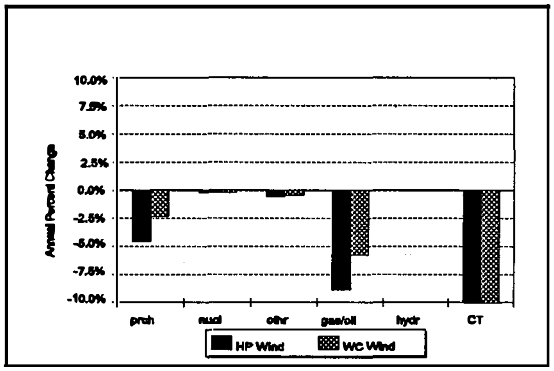

Figure 5: P+ comparison of wind/no wind cases, changes in energy

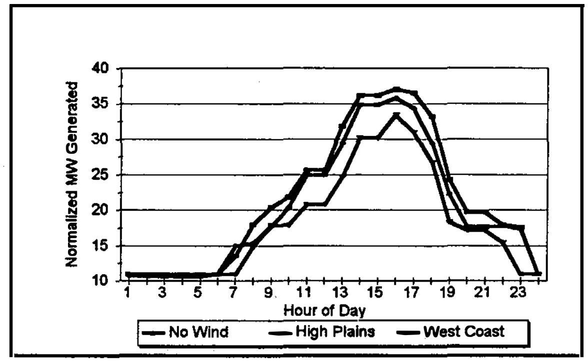

Figure 6: Utility U1 gas and oil comparison of peakday

Figure 5 summarizes the change in the use of major fuels for Utility U1 due to wind power. The change is as a percentage of annual energy generated for the no wind case. The greatest reduction is in combustion turbine (CT) and gas/oil generation. However, actual reduction in CT generation is small as its total annual generation is small. The most significant reduction is in the gas/oil units. These generators, along with the hydro units, are normally the marginal units. Whenever wind generation is available, lower cost units become the marginal units and higher cost generation is displaced. Depending on the wind regime, the annual reduction in gas/oil generation ranges from about 5% to 8% . Base-load units show extremely small changes, typically occurring during low load periods. Purchased power is reduced by approximately 1% to 3%.

Figure 6 compares the peak-day fuel reductions for the two wind regimes. The graph shows the hourly generation from the gas/oil units for no wind, HP wind, and WC wind. The WC wind results in a lower decrease in gas/oil fuel requirements during the peak day compared to HP wind.

The benefits of wind power from the HP site, as determined by P+, show a 4.1% reduction in operation cost for the year and a 49.3% reduction in annual loss of load hours (LOLH). The WC wind case showed a 2.7% reduction in costs and a 65.3% reduction in LOLH. The results show that the HP wind provides more power on the peak day, while the WC site provides greater improvement in system reliability over the full year. As for the choice between these two sites, one would need to carefully balance the reliability benefits and the economic returns offered by each site.

3.2 Elfin Results

The Elfin model produces results that are very close to those of P+. The cost benefit of wind power for the HP wind case determined by Elfin is 4.5%, and that for the WC wind is 2.6%. Due to the nonchronological nature of Elfin, daily and hourly output is not available. However, the Elfin results are consistent with P+ in that the output of the base-load units hardly changed.

In addition to the load modification method, wind power was also modelled in Elfin as a supply-side resource in two ways. The first and simplistic supply-side approach used the same time-varying wind generation schedule as the load-modification approach. When Elfin converts the chronological load and resource data to the LDC format, wind power is averaged over all hours for the given subperiod. Some temporal resolution is lost, and the statistical variance of the wind resource is significantly reduced. No attempt was made to use the unit forced-outage rate to approximate the variation of the wind resource because that information is embodied in the varying hourly wind power output.

The second supply-side representation models the wind plant as a multi-block thermal unit. Wind output is binned, according to the various inflection points on the turbine power curve, and a family of wind-speed probability distributions is calculated. One typical day is constructed for each month. This typical day is composed of one probability distribution per hour. This results in 24 probability distributions per typical day. Each of the probability distributions is calculated based on all available data for the month. The result can be interpreted as a probability distribution that describes wind power availability for each hour of a typical day per month. Therefore, we have 24 probability distributions per month, one per hour of the day, and 12 x 24 = 288 probability distributions for the full year.

Figure 7 shows the percentage reduction in annual operating costs resulting from adding wind generation to Utility U1. The first three bars show the results from the Elfin time-varying capacity (CP), load modification (LM), and multi-block probabilistic (PR) methods. The last bar shows the P+ results.

The three Elfin cases produce similar results with cost savings (benefits) of around 4.5%. Conceptually, the Elfin LM case most closely parallels the P+ case because it offers the greatest degree of temporal resolution, and deducts wind power before performing the commitment and dispatch algorithm. Comparing these two cases, we see the P+ result shows a slightly smaller benefit than the Elfin LM case: 4.1% versus 4.5%.

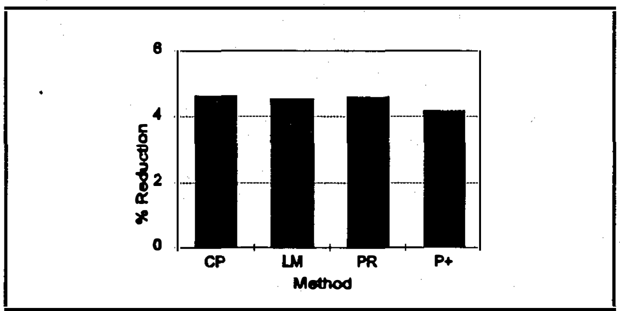

Figure 7: Cost reduction from base case, High Plains wind, Utility U1

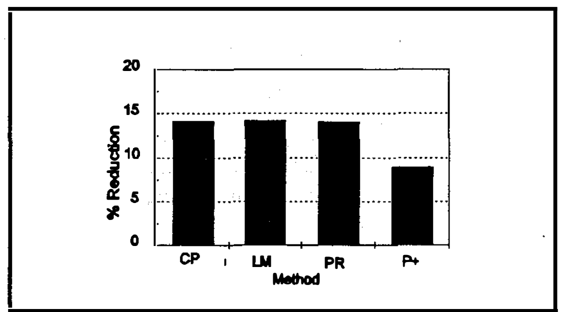

Figure 8: Change in fuel from gas and oil units, High Plains wind, Utility U1

3.3 Comparison of Model Results

The P+ model shows a slightly lower cost benefit for wind power in the HP wind case than does the Elfin model. This implies that the change in fossil fuel consumption calculated by P+ is less than that of Elfin. Since the base-load units and combustion turbines experience an insignificant absolute change, that leaves only the load-following gas/oil generators to compensate for the addition of wind power. Figure 8 illustrates the change in this group of generators. Indeed, the P+ results indicate a smaller change in the gas/oil usage than in the Elfin cases.

Two hypotheses may explain the divergence in results between the two models. First, we have not provided P+ with ramping rates. Therefore, start-up costs for slow-start units are not considered. The P+ output indicates that the number of starts and stops of gas/oil units is significant, approximately 700 unit startups for the year are eliminated with both the HP and WC wind for utility U1. Including ramp rates and other start-up information would cause P+ to result in a larger benefit for wind power.

The second hypothesis is that the constraints on unit commitment may be somewhat simplified in the Elfin model compared to the P+ model. For example, Elfin performs unit commitment for each subperiod and the units are assumed to be committed for the whole subperiod. The P+ model commits or de-commits each unit whenever possible, subject to its up and down time constraints. The commitment priority list used by Elfin differs significantly from the dynamic commitment order list used by P+. Differences in commitment will lead to dispatch differences, resulting in the different fuel consumption and cost estimates, and hence different wind benefits.

The discussion thus far has revolved around Utility U1. The analysis and model runs also were performed for Utility U2, and the results were not significantly different [8]. Base-load generation was not displaced by wind power (except for extreme low-load periods), and most of the energy displacement occurred with the gas/oil units. The calculated benefit of wind power between the two models was again very close when the load-modification method of Elfin was compared to the P+ results.

4.0 CONCLUSIONS

This study focused on the methods of incorporating wind power into production cost models. Within the LDC model framework, it appears that differences between methods may not be significant when the full year of wind data is used, although the results are sensitive to the wind regime. Furthermore, total production cost obtained by chronological model does not appear to differ significantly from that obtained by the LDC model, even though the chronological commitment appears to be more realistic and more accurate. Chronological models provide the capability of answering important questions about wind resources that are difficult or impossible to address with LDC models. Given the recent dramatic increases in computing power, the chronological models have overcome their chief disadvantage of significant performance penalties so that greater accuracy can be achieved in production cost modeling.

The value of wind power to a utility depends on the marginal generating unit(s) displaced. The marginal cost is sensitive to several factors, most importantly its definition. Because the production cost model calculates expected output from the various generators, the generators expected to be on the margin depend on their forced outage rates and dispatch order, thereby affecting the marginal cost. Furthermore, marginal cost is heavily influenced by the load level; therefore, the value of a wind plant is sensitive to the load level.

Modeling methods for intermittent resources in both chronological and LDC algorithms are still in need of some improvement. Although the load-modification method is probably the superior method because it preserves the time-varying capability of wind, neither the LDC nor the chronological modeling framework allow for both simple specifications of this time-varying feature and the uncertainty associated with intermittent generation. A possible caveat to this statement is that a promising method that was not investigated here, the branching technique in P+, may provide this capability. Future work should address this issue, expand on the probabilistic supply-side approach with both LDC and chronological models, and explore the use of Monte Carlo and related methods to capture the variance that is so important to quantify.

5.0 REFERENCES

[1] H. Baleriaux, E. Jamoulle and L. de Guerchetin, "Simulation de l'explotation d'un parc de machines thermiques de production d'electricite couple a des station de pompage", Revue E, Vol. 5, No. 7, pp. 3-24, 1967.

[2] R.R. Booth, "Power System Simulation Model based on Probability Analysis," IEEE Trans. on PAS, Vol. 91, No. 1, pp. 62-69, Jan/Feb 1972.

[3] C. Marnay and T. Strauss, *Chronological Model Comparison*, California Public Utility Commission Report, 1989.

[4] C. K. Pang, K. Lee and G. Tam, "Investigation on the Treatment of Generating Unit Outages in Probabilistic Production Simulation," Third International Conference on Probabilistic Methods Applied to Electric Power Systems, London, UK, 3-5 July 1991.

[5] M. Milligan and A. Miller, "The Value of Wind power: An Investigation Using a Qualified Production Cost Model," Wind power '93 Proceedings, San Francisco, California, 12-16 July 1993.

[6] M. Milligan, A. Miller, and F. Chapman, "Estimating the Economic Value of Wind Forecasting to Utilities," Wind power '95 Proceedings, Washington, DC, 27-30 March 1995.

[7] W. Sandusky, J. Buck, D. Renné, D. Hadley, O. Abbey, S. Bradymire and J. Gregory, *Candidate Wind Turbine Generator Site Cumulative Meteorological Data Summary and Data for January 1982 Through September 1982*, Report PNL-4663, Pacific Northwest Laboratory, Richland, Washington, 1983.

[8] M. Milligan, *Alternative Wind Power Modeling Methods Using Chronological and Load Duration Curve Production Cost Models*, Report NREL/TP-441-8171, National Renewable Energy Laboratory, Golden, Colorado, April 1996.

6.0 ACKNOWLEDGMENT

The authors would like to thank the National Renewable Energy Laboratory for the funding and support of the project, the results of which are presented in this paper.

7.0 BIOGRAPHIES

Michael Milligan is a consultant for the National Renewable Energy Laboratory, National Wind Technology Center, specializing in utility analysis and modeling. He has worked for many years in the electric utility industry, performing resource analysis, rate studies, and load forecasting. In addition to his work at NREL, he teaches computer science at a college near Denver. He has a B.A. in Mathematics from Albion College, and an M.A. and Ph.D. in Economics from the University of Colorado.

Chok Khon Pang is a founder and Principal Consultant of P Plus Corporation. He has extensive experience working in electric utility resource planning, scheduling and operation. Dr. Pang has also developed computer models on production costing, generation expansion, generation reliability, unit commitment, resource scheduling, and financial analysis. He obtained a B.E. (Honours) in Electrical Engineering and an M.Sc. in Physics from the University of Malaya, and an M.S.E.E. and Ph.D. in Electric Power Engineering from Purdue University.

Granger Tam is a founder and Senior Consultant of P Plus Corporation. He has developed computer models for power plant simulation and production costing. He is also experienced on a variety of commercial PC software, including Visual FOXPRO, ACCESS, Visual Basic, and Clipper. He obtained a B.S.M.E. from the Southeastern Massachusetts University and an M.S. in Mechanical Engineering from the Georgia Institute of Technologies.