NREL/CP-500-25793

Some Analyses of Energy Production from the NWTC Variable Speed Test Bed

Palmer W. Carlin and Lee Jay Fingersh
National Wind Technology Center

Presented at
AIAA/ASME Wind Energy Symposium
Reno, Nevada
January 11-14, 1999

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

Work performed under task number WE901420

November 1998

NOTICE

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

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

Available to the public from:
National Technical Information Service (NTIS)
U.S. Department of Commerce
5285 Port Royal Road
Springfield, VA 22161
703-605-6000 or 800-553-6847
or
DOE Information Bridge
http://www.doe.gov/bridge/home.html

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

AIAA-99-0056

Some Analyses of Energy Production from the NWTC Variable Speed Test Bed

Palmer W. Carlin, Senior Electrical Engineer
and
Lee Jay Fingersh, Staff Electrical Engineer
National Renewable Energy Laboratory
Golden, Colorado 80401

Abstract

Using the NREL 20-kilowatt, direct-drive, variable-speed wind turbine test bed, we used to collected data to examine differences between constant-speed and variable-speed wind turbine operation. The assigned goal of the variable-speed control algorithm we used was to adjust continuously the turbine rotor speed so that its tip-speed ratio (TSR) stays as close as possible to a precalculated optimum value. To examine the success of the variable-speed control, histograms of tip-speed ratio root-mean-square (rms) error for each data set were calculated and plotted. The resulting histograms of both constant and variable-speed control algorithms were compared. The results validated the expected conclusion that lower rms TSR errors are associated with higher measured wind turbine power coefficients. A second comparison was made between variable-speed histograms and synthesized constant-speed histograms for the same wind. The decrease in rms TSR error of the variable-speed case was used to quantify the observed improvement in TSR tracking brought by the variable-speed control.

Introduction

The motivation for using a mixture of variable-speed and constant-speed control algorithms at the National Renewable Energy Laboratory (NREL) variable-speed test bed was to accumulate data on which to base quantitative comparisons of the two modes of wind machine operation. The 20 kilowatt size was large enough to give confidence in scaling selected algorithms to utility-sized machines, but small enough that untried algorithms did not cause concern about possible structural damage.

To illustrate one difference between constant- and variable-speed operation of wind machines, consider the following "armchair" experiment. Suppose a variable-speed and a similar constant-speed wind machine were installed at a test site where the wind probability density spectrum consisted of a single line (with perhaps a second line at zero speed). That is to say, when the wind does blow, it blows at only one speed. Assuming that the constant-speed machine was optimized for the site, then on the basis of energy collection alone, one could not distinguish between those machines---except for the additional losses in the variable-machine's power electronics.

This statement implies that the alleged energy collection advantage of variable-speed operation will become apparent only as the wind spectrum becomes broader. If a variable-speed control algorithm is to be effective, one should be able to observe improved variable-speed performance as variation of the wind speed becomes wider.

The purpose of this paper is to present data supporting two assertions concerning mechanical energy captured at the variable-speed test bed rotor hub (not electrical output energy):

1. The variable-speed algorithms tested on ten-minute data sets produce, on average, higher power coefficients than those from constant-speed algorithms that are not optimized for these ten-minute wind samples.

2. The variable-speed algorithms tested have, in spite of high rotor inertia, succeeded in holding the average tip-speed ratio nearer the optimal target value than a constant-speed algorithm would even if the latter were optimized for the same wind sample.

Data for the second assertion show that although the enhancement is present, there is much room for

*This paper is a declared work of the U.S. Government and is not subject to copyright protection in the United States.

improvement in variable-speed algorithms regarding the handling of the rotor inertia and its stored kinetic energy.

Wind Probability Densities

Many wind machine installation sites, including the National Wind Technology Center (NWTC), can make use of at least two wind speed probability density functions. Wind energy literature[1] commonly uses the Weibull density function, which has proved to be quite useful for estimations of annual energy production. However, at the same locations, the probability density for records of less than an hour is much better approximated by a Gaussian function centered around an average wind speed. This implies that the sum of all the short records at a location will converge to a Weibull density at the end of a year. Thus since the constant-speed wind machine designer for a real installation must use the long-term function, his machine is likely to be far from optimum for some short-term wind tests.

Let us repeat the hypothetical contest described in the first paragraph in a more normal wind regime that has many ten-minute wind samples with various average values. We would expect a power collection advantage to lie with the variable-speed machine. That is, the variable-speed machine can change its rotational speed so that its optimal performance is at the center of the Gaussian wind sample. On the other hand, if the constant-speed machine happens not to be optimized for this average wind speed, it will spend most of its time operating either above or below its optimal tip-speed ratio, thus at a lower average power coefficient, and will collect less energy.

The Test Bed

The machine[2] used in these tests is a modified Grumman Windstream 33 equipped with constant chord, untwisted SERI S809 blades which form a 10-meter rotor disc. Optimal tip-speed ratio is near 7 when the blades are pitched to 3°. This rotor is loaded with a 20-kilowatt, direct-drive, permanent magnet generator. The generator, in turn, is loaded with a computer-controlled resistor bank with 256 equal steps in power. This desk top computer allows the implementation of both variable-and constant-speed algorithms by selecting appropriate load resistor combinations for the generator. Each selected load maps into a known torque at the rotor hub.

Algorithms

The constant-speed algorithm was a straightforward proportional integral control loop that continuously adjusted the generator load resistor combination to maintain the set point speed. The variable-speed algorithm used a previously calculated look-up table built around the commonly used omega-squared algorithm.[2] This algorithm causes the computer to enter the table with the present rotor speed and extract for this speed the expected power output from the rotor at maximum power coefficient. It then connects an appropriate combination of load resistors to load the generator to this power and repeats the cycle. If this load is not balanced by the wind, the power excess or deficiency will continuously drive the rotor speed toward the optimal tip-speed ratio and the system is stable. This process occurs for wind speeds from cut-in up to rated machine power. Of course, stronger winds are prevented from overloading the generator by a control transition to constant-speed at rated power value. The control system both holds constant generator torque and uses varying blade pitch to hold constant rotor speed.

Power Coefficient Determination

Experienced wind test engineers are well aware of the data scatter problems in the determination of power curves and the necessity for techniques such as binning. For example, one can calculate instantaneous variable-speed power coefficients greater than 2.5, especially when using variable-speed algorithms. Although this exceeds the Betz limit by several times, it is nevertheless correct during wind lulls in that the hub power at that instant is the sum of aerodynamic power from wind inflow and the flywheel-effect power from the deceleration of the turbine rotor. Similarly, low coefficients are measured during times of rotor acceleration when the rotor is accumulating kinetic energy.

For the purposes of this report a single number was calculated for power coefficient for the whole ten-minute data set as follows. The total instantaneous power impinging on the system was assumed to be the usual $\frac{1}{2}\rho A v_w^3$ or air density times rotor disc area times wind speed cubed.[1] The wind being taken from a single anemometer at hub height forces the implicit assumption that the total inflow across the entire disc is parallel to the rotor axis, is uniform across the rotor disc, and is instantaneously equal to the hub anemometer reading. The instantaneous power at each of the 6,000 data points in a data set was calculated, multiplied by 0.1 second to get average energy collected during that sampling period, and summed over the data set. The resulting sum was taken to be the total number of joules of energy input to the wind machine system over the ten-minute test period. Output energy was similarly calculated at each data point by finding the product of instantaneous rotor torque, angular rotor speed, and the time increment and then summing. Finally, the rotor speeds at the first and last data points were used to calculate the initial and final kinetic energies of the mechanical system. A net gain in kinetic energy represents wind energy which was captured during the ten-minute test but was not exported through the hub. Measured output energy was augmented by this small amount. Similarly, net loss was deducted from the hub output. The final number was taken to be the mechanical energy that had been extracted from the rotor hub during the ten-minute data set. The quotient of these two numbers was taken as the average power coefficient for that data set.

The Data Set

Minor variations of the above mentioned variable-speed algorithm and constant-speed algorithms ranging from 60 rpm up to 100 rpm in multiples of 10 rpm were tested in various sequences. A standard data set length of ten minutes was used. Twenty-one channels of data were collected at ten samples per second, and seven channels of blade loads and pitch were collected at 40 samples per second.

Because this paper is centered on testing variable-speed algorithms, it was essential to select only data samples consisting of pure region two operation. That is, we only analyzed ten-minute records in which the blade pitch remained constant at 3° showing power limiting had not been reached. Further, the minimum speed must exceed 49 rpm, since this speed implies that the control system was still supplying torque to the rotor.

Analysis

Although in these tests we designed the variable-speed control algorithms around the concept of trying to hold a target tip-speed ratio, the resulting data yielded insight into the intrinsic weakness of constant-speed algorithms. As mentioned above, annual wind probability densities such as the Weibull show winds over quite a wide range of speeds. However, real ten-minute records seldom show winds outside of a fairly narrow Gaussian range. This means it is a gamble whether the combination of the fixed-rotor speed and the local wind will yield a tip-speed ratio anywhere near to optimum during a particular ten-minute period.

Although viewing a plot of the routine time history of a ten-minute data set, such as Figure 1, is important as a "sanity check" to see any obvious flaws in the data, certain features are best presented using other graphical representations. In this paper histograms of tip-speed ratio have been employed. Two useful numbers can be visualized from such plots. One quantity is the departure of the center of a histogram from the target or optimal TSR for the machine. This can be quantified by calculating the RMS value of the departure from the target TSR of each data point in the ten-minute data set.

Secondly, the breadth of the histogram can be observed and is quantified by its standard deviation which is calculated in each case. For variable-speed machines this number measures the degree of success of the control system in maintaining a constant tip-speed ratio. A mathematically perfect control system would have a single line on the TSR histogram. That is, rotor speed would always be exactly proportional to wind speed.

For convenience, a variable-speed algorithm can be thought of as being similar to a self optimizing, constant-speed wind machine. Two ten-minute data sets help to illustrate this statement. Consider the two time histories shown in Figures 1 and 2. Their general appearances are similar, and their primary difference is in average wind speed. Note, however, the differences in the TSR histograms. Figure 3 shows that the lower wind of Figure 1 caused the TSR histogram to fall nearly exactly astride the optimal TSR, so that the resulting power coefficient is an acceptable 0.391. On the other hand, in Figure 4, we see the wind is much too strong for the fixed 70 rpm so the resulting small TSR values in the histogram show that the rotor is either stalled or nearly stalled most of the time. The ten-minute power coefficient in this case is a meager 0.159, and the rms value of departure from optimal TSR is almost 3 TSR units, compared to 0.9 units for the previous case.

We carried out the same calculations for most of the constant-speed data sets including speeds of 60, 70, 80, and 90 rpm. The results are shown in Figure 5. Note the uniform drop of power coefficient as the rms TSR error grows even though the ten-minute data points represent a mixture of different winds, different constant rotor speeds, and winds either too high or too low for the optimum TSR. These data validate the common assumption that fixed-speed wind machines should be supplied with gear boxes that place their best power coefficient on the most productive wind speed at that site.

A similar plot of the variable-speed algorithm data sets is shown in Figure 6. Although the major peak of the TSR histograms in each of these cases is close to the optimum TSR of 7, the fact that there are standard deviations of order one shows that these peaks could still be made narrower with an improved control system. The cases of larger TSR error are typically caused by appreciable periods of low wind where the rotor is slowing. These samples create a large tail on the histogram out to the right such as shown in Figure 7.

For comparison of variable-and constant-speed control, Figures 5 and 6 are merged into Figure 8.

Evaluating a Variable-Speed Controller

The other major information that was extracted from the TSR histograms follows from the definition of TSR, which is $\omega R/V_w$ where omega is angular rotor speed, R is the radius of the rotor swept area, and V_w is the wind speed. Thus all wind machines operating at the same constant-speed, having the same size rotor disk, and subjected to the same wind sample will have identical TSR histograms.

Taking advantage of this fact allows one to learn more from a variable-speed data set. Recalculating the TSR with a constant number for the rotor speed will yield the histogram which would have resulted had the machine been operating at that constant-speed. The position of the histogram along the TSR axis will be set by the value chosen for the assumed rotor speed. By inserting the average rotor speed of the variable-speed data set, the resulting constant-speed histogram will fall nearly on top of the experimental variable-speed histogram. It is now possible to supply evidence on the other possible property of a variable-speed controller, namely that such a controller can extract more energy than a constant-speed controller, even though the latter may be optimized for that site. Although this approach does not answer this question directly, it can provide evidence about whether the variable-speed controller holds the TSR more closely to the target optimum than the optimized constant-speed controller.

An example of two TSR histograms is shown in Figure 9. The higher and narrower peak of the variable-speed histogram shows that that controller had, on average, influenced rotor speed sufficiently to follow the wind's major changes and thus hold the TSR closer to the target value.

This tendency was quantified by calculating the standard deviation of both histograms with the expectation that the variable-speed record would be narrower than that of the constant-speed controller. About sixteen other variable-speed data sets were processed and the results are plotted in Figure 10. The two standard deviations were plotted along the two axes. Data sets whose points lie near the line y = x (labeled "No Diff") show that the variable-speed controller did not improve the machine TSR tracking very much. Note that these points tend to have low values of standard deviation, implying that the histograms are already somewhat narrow. On the other hand, many of the points show that the variable-controller reduces the histogram scatter to between 60% and 80% of the constant-speed value, thereby implying an operating locus in the TSR-variable-speed plane that remains nearer the target TSR.

Summary

The original and continuing purpose of the NREL direct-drive, variable-speed experiment is to gather experimental data to help understand the mechanical transformation of wind kinetic energy to mechanical energy at a wind machine rotor hub. Efficiency of other power handling components such as gear boxes or power electronics were not considered. Most of the data we collected was analyzed in this paper from the standpoint of wind machine tip-speed ratio. Although power coefficient is of greater final importance, we have assumed that smaller excursions of TSR away from a machine's optimum TSR are a prerequisite for higher power coefficients.

We draw two conclusions from the tests. The first is that evidence supports the conclusion that, for a constant-speed machine, the farther its TSR is from its optimal value, the worse its average power coefficient will be. This implies that a constant-speed machine should be optimized for the wind regime into which it will be installed. Variable-speed machines for a broad range of speeds can be thought of as being able, on average, to optimize themselves to the existing wind.

The second conclusion is that one measure of performance of a variable-speed wind controller can be quantified by comparing the breadth of its tip-speed ratio histogram with its implied constant-speed histogram. It will be interesting to see how narrow such a histogram can be made. However for the practical wind machine designer such histograms imply much higher rotor accelerations which could mean much higher blade root loads and much higher fatigue.

References

[1] D.M. Eggleston, F.S. Stoddard, Wind Turbine Engineering Design, p 67, Van Nostsrand Reinhold Company, New York, 1987.

[2] L.J. Fingersh, P.W. Carlin, "Results from the NREL Variable-Speed Testbed," 1998 ASME Wind Energy Symposium, Reno, NV, January 12-15, 1998.

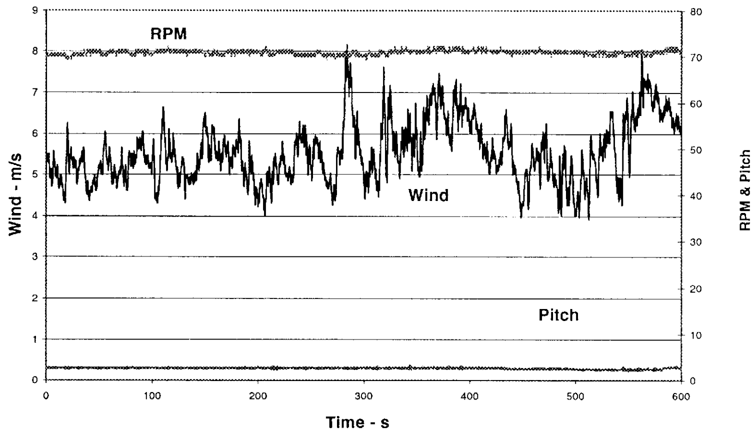

Ten minute record using constant 70 rpm algorithm

Figure 1

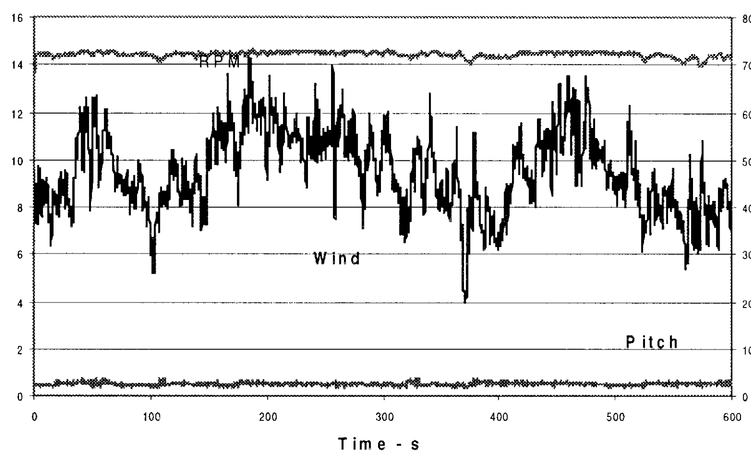

Ten minute record using 70 rpm constant speed algorithm

Figure 2

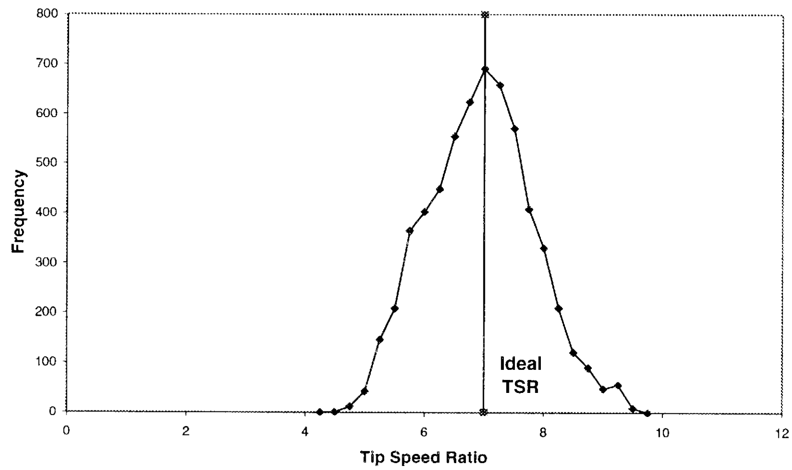

Tip-Speed Ratio Histogram for 5 m/s average wind

Figure 3

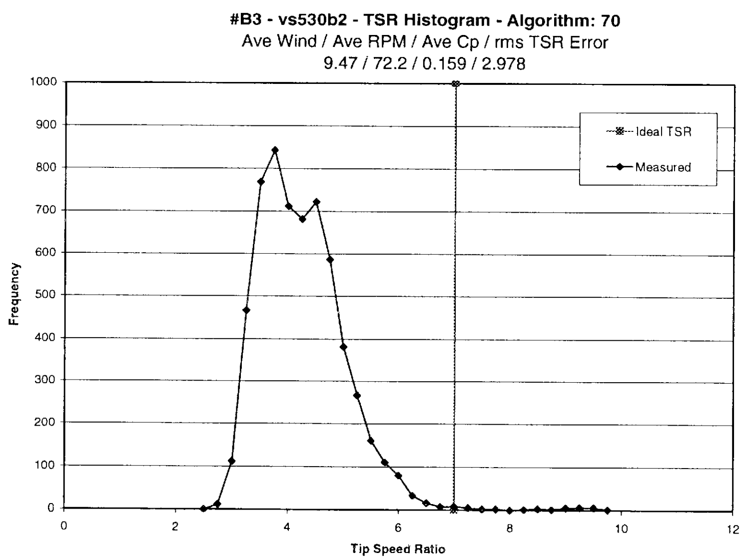

Figure 4

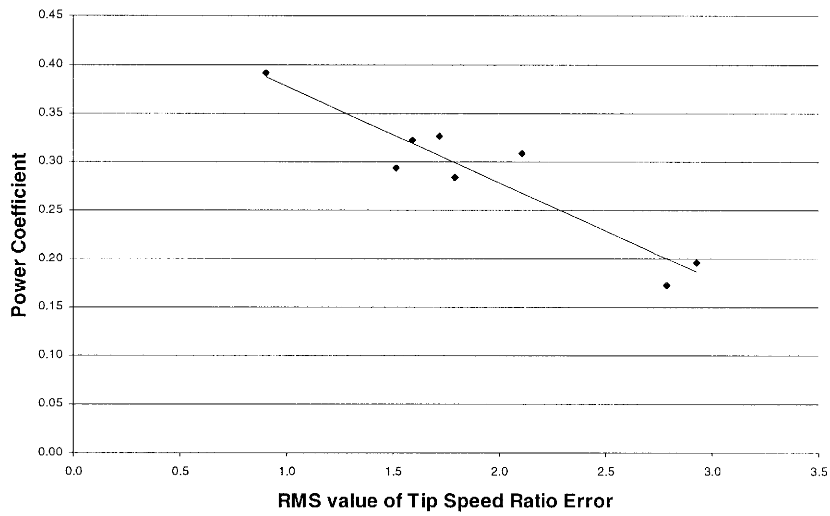

Various Constant Speed Algorithms with Trend Line

Figure 5

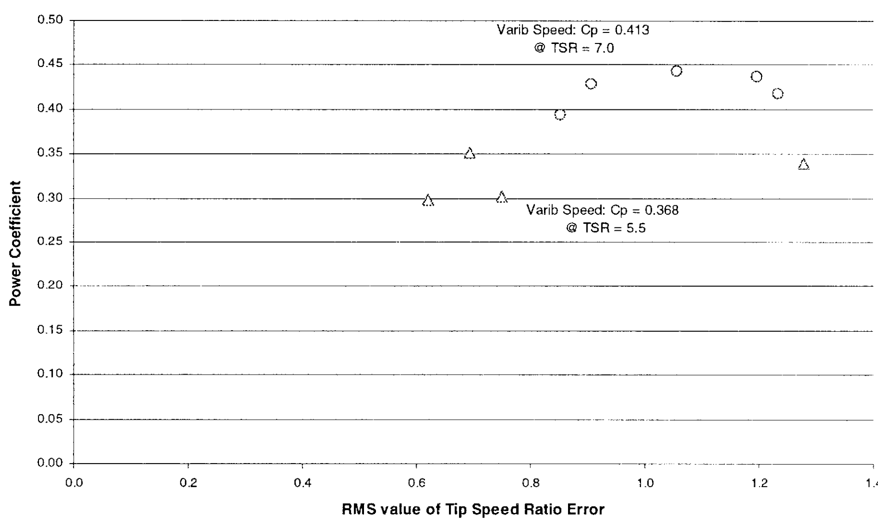

Power Coefficients for Two Varible Speed Algorithms

Figure 6

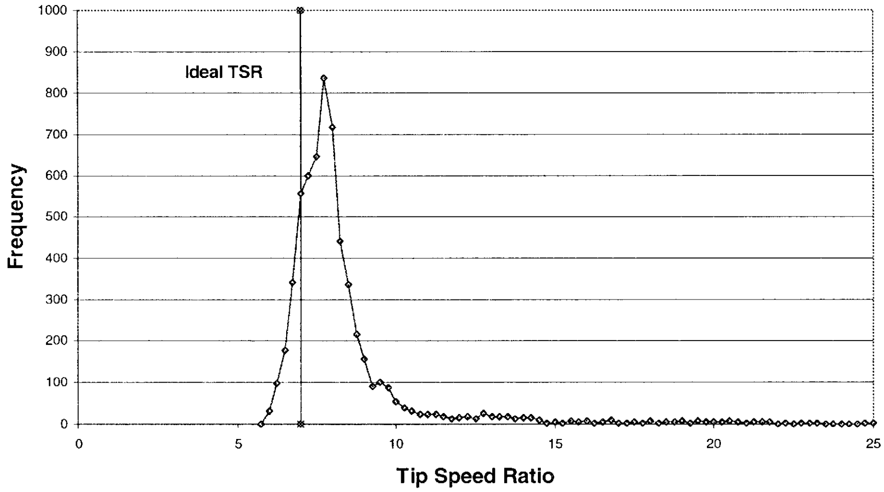

Tip-Speed Ratio Histogram for Varible-Speed in low wind

Figure 7

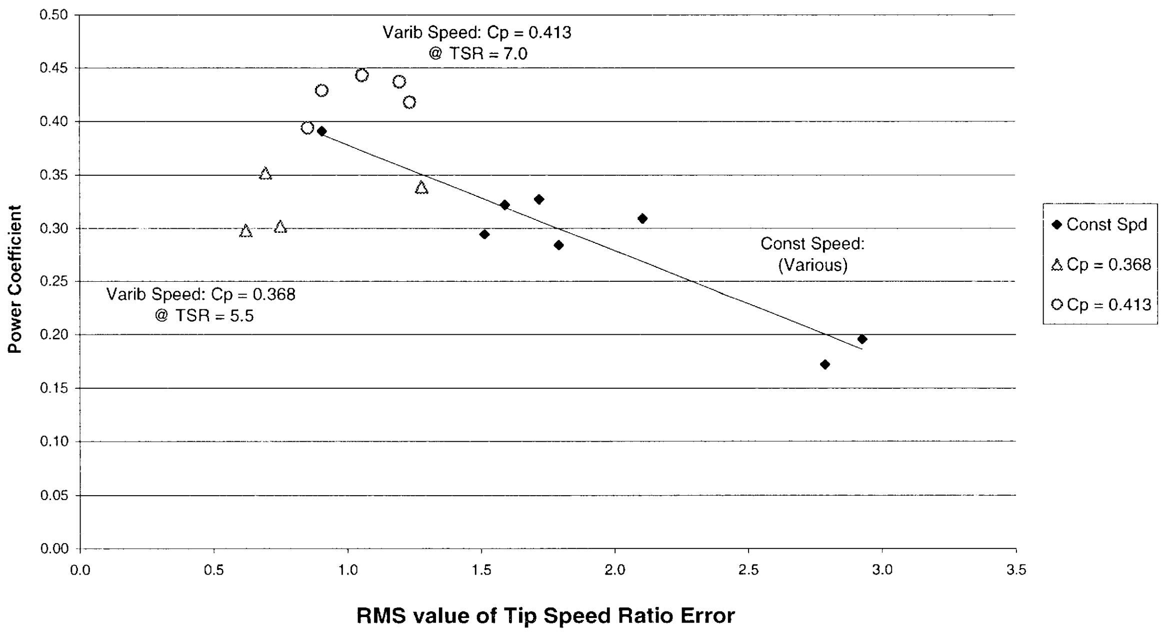

Comparison of Varible Speed and Constant Speed Algorithms

Figure 8

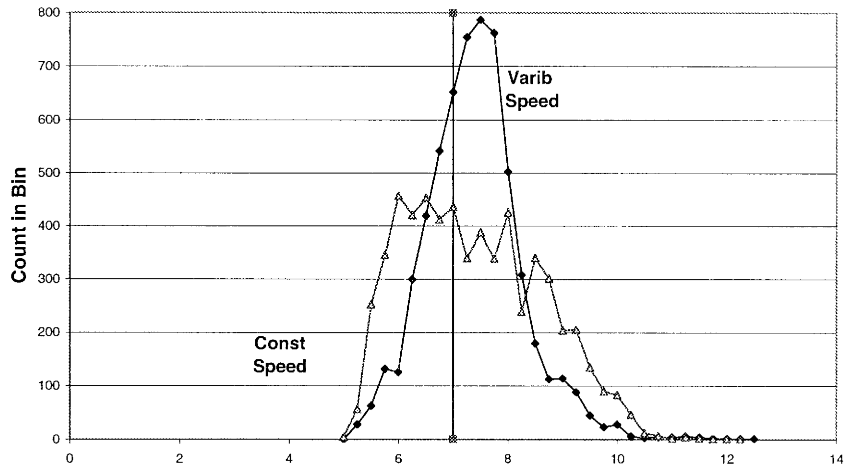

Tip Speed Ratio

Comparison of Variable- and Constant-Speed Algorithms
Using the Same Wind Sample

Figure 9

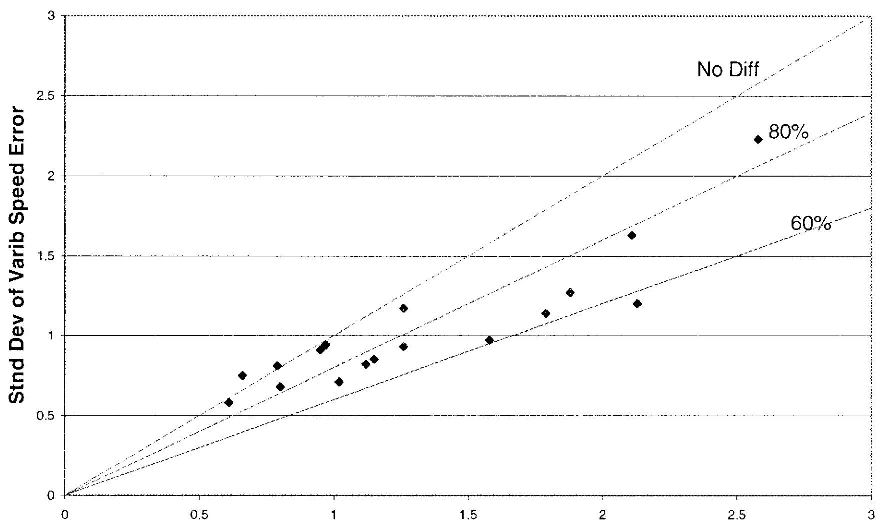

Stnd Dev of Const Speed TSR Error

Tracking Improvement Afforded by Variable Speed Algorithm

Figure 10