PREDICTION OF POWER GENERATION FROM A WIND FARM

Bill Whiten* and Tasos Tsoularis†

Abstract

Wind farms produce a variable power output depending on the wind speed. For management of power networks and for bidding for the supply of power, the future power available needs to be predicted for time intervals ahead of a few minutes to about 24 hours.

This project used data from a wind farm and three meteorological stations to determine methods and ability to predict wind speed. Analyses using regression, neural networks, and a Kalman filter were examined. Prediction using a combination of local wind measurements and meteorological data appears to give the best results.

1. Introduction

Electrical power as such is not stored and hence generation must match demand. Chemical or hydraulic energy storage and regeneration is possible on a medium scale, but is inefficient and difficult on the scale of a national grid. Wind power generation (and solar power) differs from most other large scale power generation in that the amount of energy available is set by natural conditions rather than being under manual control. This makes balancing supply and demand more difficult particularly when, as is expected in the future, a large proportion of the grid power is coming from such sources.

When wind power is used it is necessary to predict the amount of power that will be available for future periods from five minutes up to a day for both balancing load and for bidding for the supply of power.

*JKMRC, University of Queensland, Australia. E-mail W.Whiten@uq.edu.au
†Massey University at Albany, Auckland, NZ. E-mailA.D.Tsoularis@massey.ac.nz

Wind power has minimal undesirable environmental effects and is expected to become of increasing importance. However it has been found to be particularly variable and difficult to predict. Variations in wind speed occur in the time scale of minutes due to local turbulence, up to several days for typical meteorological structures, and at a scale of months for seasonal variations.

The amount of power available from wind varies considerably with wind speed. For low wind speeds the power available is related to the cube of the wind speed (the energy is proportional to velocity squared times the mass, and the available mass is proportional to wind speed giving the cubic relation). Intermediate wind speeds give constant power at the design rating of the generator, and high wind speeds require the shut down of the generator for safety reasons and hence zero power is generated.

Small amounts of wind power are absorbed into the normal variations in demand. However as the proportion of wind power on a network increases it becomes the major source of non systematic variation on the power network. This then requires both good prediction of the amount of power that will be generated and sufficient means to provide an alternative variable amount of power to compensate for the variations in wind power.

Load balancing (of electrical supply versus demand) is done on a five or ten minute basis with the allocation of generation units being done two hours ahead. Bids for the supply of power at a given price are initially made 24 hours ahead and become final two hours ahead. In the cases of wind power it is difficult to predict the amount of power that will be available and hence be able to make firm bids for supply.

This project aimed to examine the prediction of the amount of power available from a wind generations installation over the time scales of interest for load balancing and bidding for supply. As there is a direct relation between the wind speed and the power generated and details of the available on power generation did not include individual turbines or the number of turbines operating, the main emphasis was placed on the prediction of wind speed.

There are three methods available for wind speed prediction. The first uses past wind speed data taken at the site to estimate future wind speeds. This data includes speeds measured right up to the time the prediction is made. The second uses meteorological forecasts made from data collected at a range of widely separated sites. This data is used as the basis for a fluid dynamic simulation to predict future weather including wind velocities. A third method is to use measurements taken up wind of the generators (hence in several different directions) to provide estimates of wind changes in advance of the change. No suitable data was available to investigate this third method so this report concentrates on the first two methods.

The Tararua wind farm supplied data on wind speed and power generation from their wind farm site. The industry sponsor, Transpower, arranged for forecast data from three sites (Masterton, Palmerston North, and Wellington) to be made available for the MISG.

There is a considerable literature on the forecasting of wind speed, most of which reports small improvements (typically 5% to 20%) over the simple persistence method (i.e. the future wind speed is predicted to be the current wind speed!). The European project ANEMOS has produced a recent and comprehensive report on the state of the art in short term prediction of wind power (Giebel 2003). This and the references therein provides an excellent introduction to forecasting wind speed.

2. Regression analysis of meteorological data

Meteorological data for Masterton, Palmerston North, and Wellington were supplied for the period 1st October 2003 to 16th January 2004. This data contains 864 measured instantaneous wind speed and direction measurements at 3 hourly intervals, and 432 sets of forecast data. The forecasts made every six hours, are for 3 to 24 hours ahead in 3 hour intervals, and give wind speed and direction from an initial forecast and from a statistically corrected forecast.

The first columns of table 1 give the standard deviation of the errors in prediction for the persistence method (i.e. prediction is simply the current wind speed), the meteorological initial forecast, and the corrected forecast. In the case of Masterton this data is plotted in figure 1. The corrected forecast is more accurate than the initial forecast, while the persistence method becomes less accurate as the prediction time increases. The predictions at 24 hours ahead are more accurate than the preceding predictions as there is a tendency for the wind speed to be similar at the same time of day. Figure 2 plots the persistence predictions against the measured values. It can be seen there is little correlation between the observed and predicted wind speed after the first few hours.

Masterton: Fitted standard errors (daily average fit 4.7)						
Time	Persistence	Forecast	Corrected	Regression	Reg+fore	Reg+corr
3	4.6	4.9	4.3	3.8	3.5	3.4
6	5.4	5.4	4.9	4.1	3.7	3.7
9	6.4	6.0	5.4	4.3	3.9	3.8
12	6.6	5.8	5.6	4.5	4.1	4.0
15	6.6	5.5	5.4	4.5	4.0	3.9
18	6.2	5.7	5.4	4.5	4.0	3.9
21	5.9	5.5	5.1	4.5	4.1	4.0
24	5.9	5.4	5.2	4.7	4.0	3.9
Fitted standard errors (daily average fit 5.0)						
Time	Persistence	Forecast	Corrected	Regression	Reg+fore	Reg+corr
3	4.1	4.9	4.3	3.2	3.0	2.9
6	5.3	5.5	5.2	3.8	3.4	3.2
9	6.7	6.4	6.1	4.4	3.7	3.6
12	6.9	6.6	6.4	4.4	3.8	3.6
15	6.8	6.4	6.2	4.7	4.0	3.8
18	6.6	6.2	5.9	4.7	3.9	3.7
21	6.0	5.9	5.3	4.8	4.3	4.2
24	6.0	5.7	5.1	4.8	4.0	3.8
Fitted standard errors (daily average fit 5.0)						
Time	Persistence	Forecast	Corrected	Regression	Reg+fore	Reg+corr
3	4.3	7.2	5.2	4.0	3.7	3.7
6	6.0	8.1	6.4	5.0	4.3	4.5
9	7.1	8.7	6.9	5.7	5.0	5.1
12	8.0	9.2	7.7	6.0	5.1	5.4
15	8.2	9.4	8.0	6.2	5.4	5.5
18	8.5	9.3	8.1	6.2	5.2	5.5
21	8.5	9.3	7.8	6.3	5.5	5.7
24	8.9	9.3	8.0	6.4	5.3	5.6

Table 1. Fitted standard errors

Various linear regressions were examined. As there is a considerable number of data points, terms that made very small changes to the accuracy of the prediction were reported as being significant under the standard statistical tests. In the interest of robustness only terms that have both a noticeable effect on accuracy and a clear physical meaning are recommended.

A simple regression from the current wind speed and a constant showed the coefficient of the current wind speed decreased and that of the constant increased as the time ahead of the prediction increased. In other words the measured wind speed gave decreasing value to the prediction as the time ahead increased, and the prediction trended towards a simple constant.

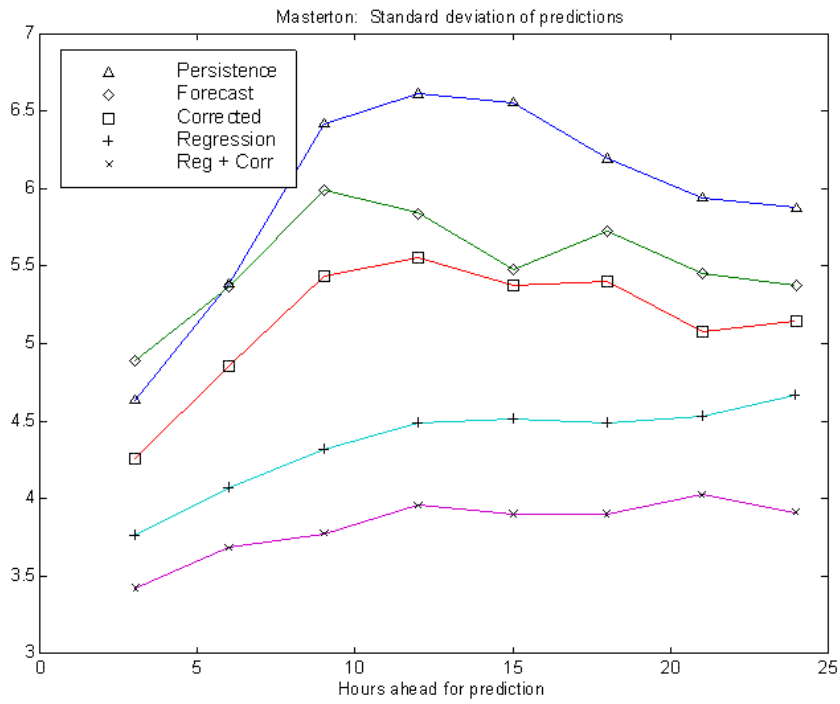

Figure 1. Standard deviations of the errors for the different types of predictions for Masterton.

A time of day variation was found (figure 3) and prediction using the average for the time of day only was quite accurate (standard deviation for Masterton: 4.7, Palmerston: 5.0, and Wellington: 6.4). From the time of day figures the expected change in wind speed, between the time at which the prediction is made, and the time for which the prediction is made, can be calculated. This value can be included as a data column in the regression, and tests showed it gave a noticeable improvement in accuracy.

Instead of the regressions tending towards a constant value as the time ahead increases, the regression was altered so that it could tend towards the average for the time of day as the prediction time increased. This term gave a small improvement in the accuracy of the predictions.

Thus the recommended regression prediction is a linear combination of the current wind speed, the average wind speed for the time of day of

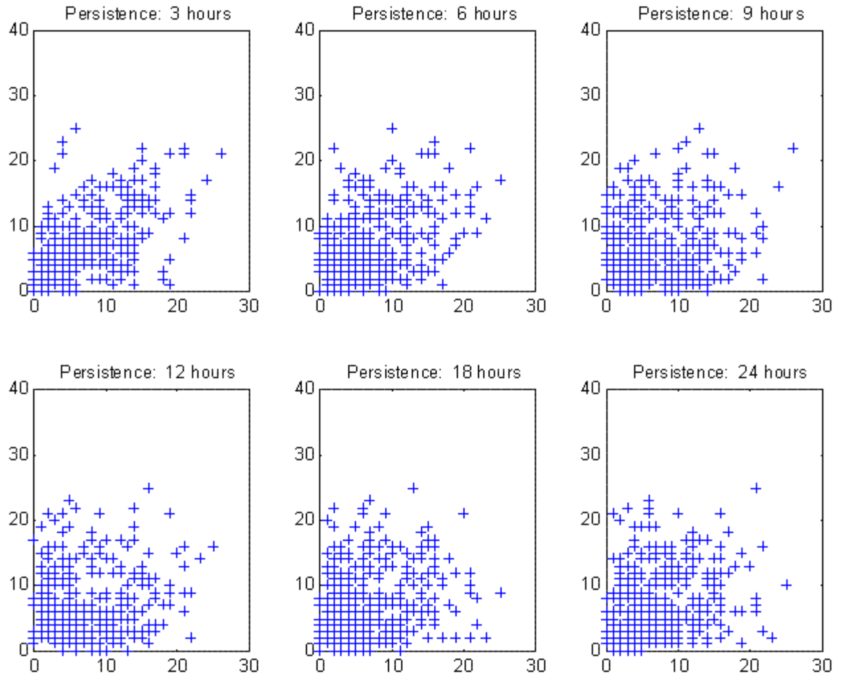

Figure 2. Persistence predictions (current wind speed) plotted against measured wind velocity for Masterton (x axis).

the prediction, and the change in average wind speed between the time of the prediction was made and the time of the prediction i.e.:

$S(t+h)$ is predicted by: $a1S(t) + a2D(t+h) + a3(D(t+h) - D(t))$

where $S(t)$ is the wind speed at time t, $D(t)$ is the average daily wind speed for time t, and h is the time ahead for which the prediction is made. The coefficients $a1$, $a2$, and $a3$ are determined by linear regression which is the minimisation of the squared difference between the predicted value and the corresponding measured value.

All terms in these regressions were highly significant. The fourth column in table 1 gives the standard deviation of the error for this regression. For a prediction 24 hours ahead the prediction is essentially the same as that of the daily average value. The coefficients ($a1$ to $a3$) for each time interval are given in table 2. All the coefficients are positive and have sufficient physical meaning to give confidence that the regression will be quite robust. As noted for the simple regression (wind speed

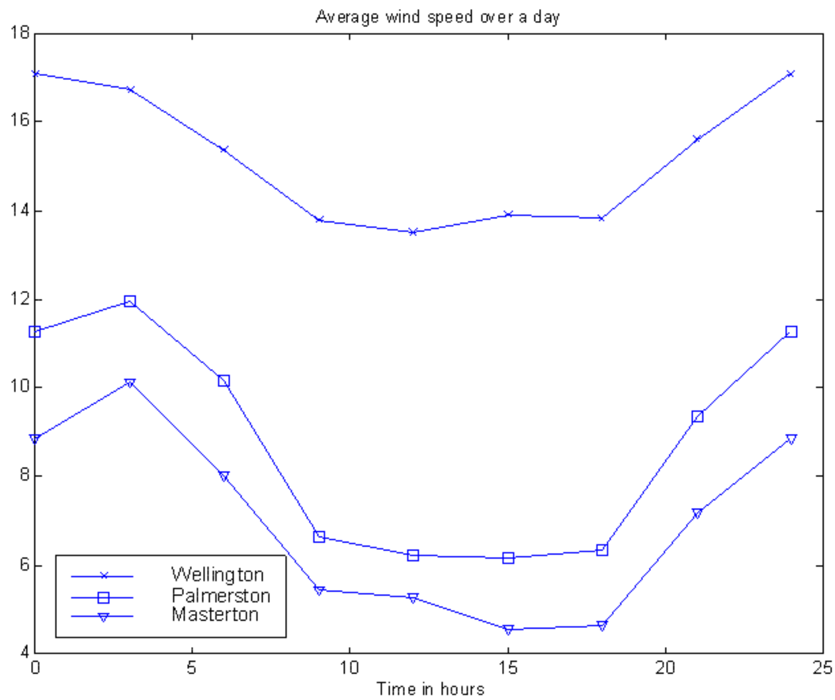

Figure 3. Variation of average wind speed with time of day.

and a constant) the coefficients change from at 3 hours a major reliance on the measured wind speed, to at 24 hours ahead a significant bias to the average wind speed term. It should be noted that the available data covers only one part of the year and the average daily wind behaviour may be different at other times of the year.

3. Inclusion of Meteorological Forecasts

The meteorological forecasts have the advantage of having significant additional data, both in time and space, and are based on physically based formula coded into a numeric prediction. However these forecasts do not exactly correspond to local measurements. By combining the forecast value and the regression prediction in a second regression we can develop formula that chooses the best combination of the two values. As the two values (regression prediction and forecast) both are predictions of the wind speed, a linear combination of them with posi-

Masterton: Regression coefficients

Time	Measured	TimeConst	Adjust
3	0.591	0.411	0.592
6	0.483	0.521	0.489
9	0.376	0.624	0.373
12	0.266	0.735	0.268
15	0.250	0.750	0.246
18	0.262	0.740	0.261
21	0.220	0.782	0.203
24	0.258	0.711	0.000

Palmerston: Regression coefficients

Time	Measured	TimeConst	Adjust
3	0.793	0.205	0.783
6	0.623	0.376	0.619
9	0.503	0.494	0.497
12	0.420	0.581	0.421
15	0.371	0.629	0.369
18	0.291	0.710	0.292
21	0.304	0.693	0.300
24	0.289	0.699	0.000

Wellington: Regression coefficients

Time	Measured	TimeConst	Adjust
3	0.794	0.206	0.794
6	0.603	0.398	0.611
9	0.453	0.546	0.446
12	0.301	0.699	0.306
15	0.231	0.769	0.221
18	0.142	0.860	0.143
21	0.121	0.880	0.105
24	0.025	0.972	0.000

Table 2. Regression coefficients for measured wind speed, average wind speed for time of day, and expected change in wind speed.

Masterton Regression + Corrected: coefficients

Time	Regression	Corrected
3	0.530	0.500
6	0.421	0.572
9	0.350	0.668
12	0.263	0.717
15	0.270	0.759
18	0.314	0.693
21	0.318	0.717
24	0.232	0.777

Palmerston Regression + Corrected: coefficients

Time	Regression	Corrected
3	0.634	0.428
6	0.390	0.670
9	0.324	0.763
12	0.206	0.881
15	0.241	0.860
18	0.153	0.939
21	0.317	0.792
24	0.083	0.996

Wellington Regression + Corrected: coefficients

Time	Regression	Corrected
3	0.689	0.359
6	0.554	0.509
9	0.462	0.625
12	0.500	0.579
15	0.479	0.606
18	0.461	0.627
21	0.452	0.637
24	0.419	0.682

Table 3. Coefficient for combining regression prediction and the corrected forecast.

tive coefficients should give a robust prediction. In particular it will still give adequate service if the forecasts are improved (although the prediction would be improved by recalculation of the regression coefficients). Table 3 gives the coefficients for this regression based on the corrected forecast. The coefficients weight towards the regression prediction at low time intervals, and toward the meteorological prediction at longer times.

Accuracies for the regression prediction, and the combined regression and corrected forecast prediction, are given in table 1 and shown for Masterton in figure 1. A plot of the measured values and the predicted values is given in figure 4 which can be compared with figure 2 for persistence only. It can be seen that the combined regression gives better predictions than simple persistence, however there is still significant scatter.

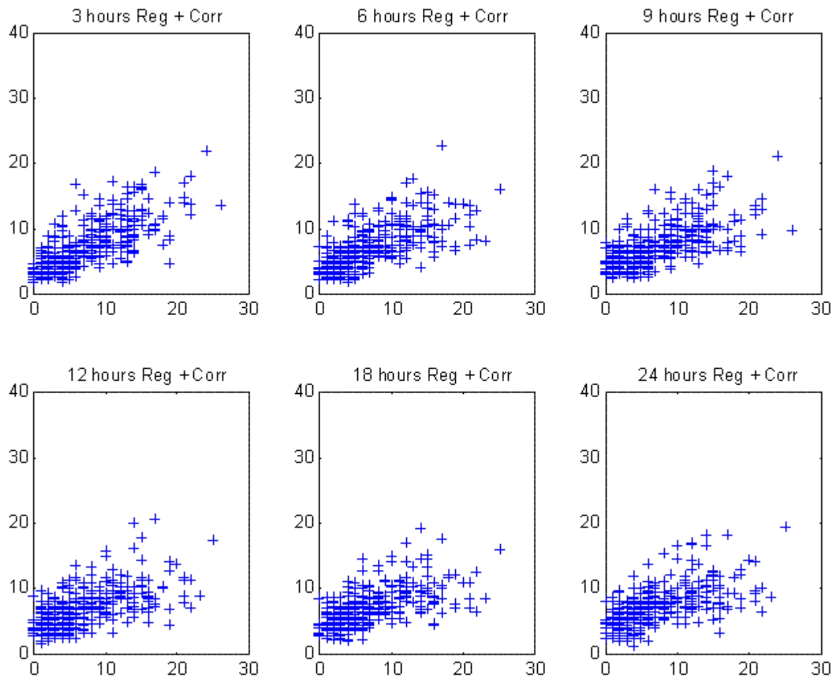

Figure 4. Predicted speed on the y axis and measured speed on x axis. Note how the prediction approaches a constant as the time ahead increases.

An error distribution for the various predictions can be calculated. Table 4 shows how the error distribution changes with prediction time for persistence predictions, and table 5 shows the error distributions for the combined regression and corrected forecast. It can be seen that quite large differences between predicted and actual wind speed can occur even for the shorter prediction times.

Comments on Linear Regression of the Meteorological Data

The best results were obtained using a combination of a local wind speed prediction and the meteorological forecasts. The local prediction dominated the short term predictions while the average values and forecast values dominated the longer term predictions.

The method of combining a local prediction with that from a meteorology forecast, can be applied to other local predictions as well as the

Masterton: Persistence: error distribution						
time	1%	5%	20%	80%	95%	99%
3	-14.5	-7.0	-3.0	3.0	8.0	13.0
6	-13.0	-9.0	-4.0	4.0	9.5	12.0
9	-15.0	-11.0	-5.0	6.0	11.0	14.0
12	-17.7	-10.5	-6.0	5.0	11.0	13.7
15	-15.0	-11.0	-5.0	5.0	11.0	15.7
18	-16.0	-11.0	-5.0	5.0	9.0	15.7
21	-16.4	-10.0	-4.0	4.0	9.0	14.0
24	-16.7	-10.0	-4.0	4.0	10.0	14.7

Palmerston: Persistence: error distribution						
time	1%	5%	20%	80%	95%	99%
3	-9.0	-7.0	-3.0	3.0	7.0	11.0
6	-12.0	-8.0	-4.0	4.0	9.0	13.0
9	-14.0	-10.4	-6.0	6.0	11.0	16.7
12	-16.0	-11.0	-6.0	6.0	11.0	15.4
15	-17.7	-11.0	-6.0	11.4	11.4	15.7
18	-15.7	-11.0	-5.0	6.0	10.0	15.0
21	-14.0	-10.0	-5.0	5.0	10.0	14.7
24	-13.0	-10.0	-5.0	5.0	10.0	14.7

Wellington: Persistence: error distribution						
time	1%	5%	20%	80%	95%	99%
3	-9.7	-7.0	-3.0	3.0	8.0	11.0
6	-13.7	-9.4	-5.0	5.0	10.4	14.4
9	-16.7	-12.0	-6.0	6.0	12.0	16.0
12	-21.0	-14.0	-6.6	7.0	13.0	17.7
15	-19.7	-14.0	-7.0	7.0	13.0	17.7
18	-23.7	-15.0	-7.0	7.0	13.0	19.4
21	-21.0	-14.0	-7.0	8.0	13.4	19.7
24	-21.7	-15.0	-8.0	7.6	14.0	20.0

Table 4. Distributions of error in prediction for persistence method.

Masterton: Regression + Corrected: error distribution

time	1%	5%	20%	80%	95%	99%
3	-7.5	-5.1	-2.7	2.2	6.3	10.0
6	-7.4	-5.1	-3.0	2.4	7.3	11.4
9	-7.8	-5.3	-3.2	2.6	6.9	10.2
12	-7.2	-5.5	-3.1	2.5	7.9	12.1
15	-7.3	-6.0	-3.2	2.7	7.0	10.8
18	-7.0	-5.2	-3.1	2.4	7.5	12.4
21	-7.5	-5.6	-3.4	2.8	7.3	11.3
24	-7.8	-5.1	-3.2	2.6	7.8	12.1

Palmerston: Regression + Corrected: error distribution

time	1%	5%	20%	80%	95%	99%
3	-6.7	-4.4	-2.5	2.1	4.6	7.3
6	-7.2	-4.9	-2.8	2.4	5.2	9.1
9	-9.0	-5.2	-3.2	2.5	6.5	9.7
12	-7.3	-5.0	-3.1	2.9	5.7	9.4
15	-7.2	-5.3	-3.4	2.9	6.7	11.3
18	-8.3	-5.3	-3.2	3.0	6.0	10.4
21	-8.7	-5.7	-3.4	3.1	7.3	11.6
24	-8.8	-5.6	-2.9	3.1	6.3	10.4

Wellinton: Regression + Corrected: error distribution

time	1%	5%	20%	80%	95%	99%
3	-8.7	-6.0	-3.1	2.9	6.1	8.5
6	-11.0	-7.7	-3.8	3.4	7.0	9.2
9	-11.8	-9.2	-4.4	4.3	7.9	10.6
12	-12.6	-9.5	-4.8	4.9	8.6	11.0
15	-12.6	-9.3	-5.1	4.6	9.3	11.6
18	-12.4	-9.8	-4.8	4.7	9.4	11.1
21	-13.0	-10.1	-5.3	4.5	9.5	12.2
24	-13.3	-9.7	-4.7	4.8	8.8	11.3

Table 5. Distributions of error in prediction for combined regression prediction and corrected forecast.

regression prediction demonstrated here, As the forecasts are based on relevant data not available from local measurements this technique is recommended to enhance predictions made using local data.

The wind speed measurements are instantaneous values rather than an average as expected from the predicted and forecast values. It is not at all clear how much of the variation is due to the measurements and hence is inherent in the comparisons. Extrapolating the graphs (in figure 1) back to zero hours indicates there may be a significant amount of variation in the measurements, If this exists and is subtracted the proportional changes in the prediction accuracies will be significantly increased.

There are several possible options that might improve the regressions to predict wind speed that have not been fully investigated. The wind direction may provide a term that improves the regression. The regressions could instead of predicting wind speed, predict the two components of the wind velocity. A superficial investigation of these did not indicate large gains.

The daily wind profile could be investigated further to determine profiles that correspond to different weather conditions and/or seasons. The available data did not allow a detailed investigation of this.

Nonlinear terms such as powers or spline functions can be easily included in the linear regressions. This was not tested, however the amount of apparently random variation present indicates the gain from this may be minimal.

Another possibility is to use log scales rather than linear scales, which should make the error distribution more symmetric, but does not allow for the zero speed values in the data.

4. Regression Analysis of Wind Farm Data

The wind farm data consists of 204721 ten minute average values of wind speed, wind direction, and power generated over the period 12/3/99 to 2/7/03 (1574 days). About 10% of the values over this period are missing from the data, and the wind speed and power values have been normalised to the range zero to one. Information on the number of windmills operating is not included, which limits the ability to predict power output.

Table 6 gives the predictions using only persistence for selected time intervals ahead. As expected the quality of the predictions reduces as the time increases, and eventually becomes less accurate than prediction by a simple constant which gives a standard deviation of 0.153. Table 7 gives the results for a simple regression using the current 10 minute

Hr:Min	S.D.	1%	5%	20%	80%	95%	99%
0:10	0.0230	-0.06	-0.04	-0.02	0.02	0.04	0.06
0:20	0.0323	-0.09	-0.05	-0.02	0.02	0.05	0.09
0:30	0.0386	-0.10	-0.06	-0.03	0.03	0.06	0.10
1:00	0.0524	-0.14	-0.08	-0.04	0.04	0.08	0.14
2:00	0.0710	-0.18	-0.12	-0.05	0.05	0.11	0.18
3:00	0.0850	-0.22	-0.14	-0.06	0.06	0.14	0.22
6:00	0.1149	-0.29	-0.19	-0.09	0.08	0.19	0.30
12:00	0.1497	-0.38	-0.24	-0.11	0.11	0.24	0.39
18:00	0.1678	-0.42	-0.28	-0.13	0.13	0.28	0.42
24:00	0.1781	-0.44	-0.30	-0.14	0.14	0.30	0.44
48:00	0.1963	-0.47	-0.33	-0.16	0.16	0.32	0.48

Table 6. Standard deviation and error distribution for persistence method at different time intervals.

Hr:Min	S.D.	1%	5%	20%	80%	95%	99%
0:10	0.0229	-0.06	-0.04	-0.02	0.02	0.04	0.06
0:20	0.0321	-0.08	-0.05	-0.02	0.02	0.05	0.09
0:30	0.0383	-0.10	-0.06	-0.03	0.03	0.06	0.11
1:00	0.0516	-0.13	-0.08	-0.04	0.04	0.09	0.14
2:00	0.0691	-0.17	-0.11	-0.05	0.05	0.12	0.19
3:00	0.0816	-0.19	-0.12	-0.06	0.06	0.14	0.22
6:00	0.1064	-0.25	-0.16	-0.08	0.08	0.18	0.28
12:00	0.1305	-0.28	-0.19	-0.11	0.10	0.24	0.35
18:00	0.1401	-0.27	-0.20	-0.12	0.12	0.26	0.37
24:00	0.1446	-0.26	-0.20	-0.12	0.12	0.27	0.38
48:00	0.1502	-0.25	-0.20	-0.13	0.13	0.28	0.39

Table 7. Standard deviation and error distribution for simple regression (current wind speed and a constant) at different time intervals

average and a constant and table 8 gives the coefficient values for this regression. Again all the coefficients are highly significant (t-test values above 30). Similar to the meteorological data the measured wind speed becomes less important and the prediction tends towards a constant as the time increases. The use of the previous two ten minute average wind speed gave almost zero change in accuracy and weighted the regression equations heavily towards the more recent ten minutes.

The average daily variation for the wind farm data (figure 5, note the expanded scale compared with figure 2) is not as great as for the three sets of meteorological data, possibly due to this data covering a full year. Regression using the current wind speed, daily average for the time for which the prediction is made, and the average change in wind speed between the two times (as for the meteorological data) gave

Hr:Min	Coef	Constant
0:10	0.989	0.003
0:20	0.978	0.006
0:30	0.968	0.009
1:00	0.941	0.017
2:00	0.892	0.031
3:00	0.845	0.044
6:00	0.717	0.081
12:00	0.518	0.138
18:00	0.393	0.173
24:00	0.315	0.195
48:00	0.173	0.236

Table 8. Regression coefficients for current wind speed and a constant at different time intervals.

Hr:Min	S.D.	1%	5%	20%	80%	95%	99%
0:10	0.0229	0.03	0.02	0.01	0.01	0.03	0.04
0:20	0.0321	0.06	0.04	0.02	0.02	0.04	0.06
0:30	0.0383	0.08	0.05	0.02	0.02	0.05	0.09
1:00	0.0515	0.10	0.06	0.03	0.03	0.06	0.11
2:00	0.0688	0.13	0.08	0.04	0.04	0.09	0.14
3:00	0.0811	0.17	0.11	0.05	0.05	0.12	0.19
6:00	0.1055	0.19	0.12	0.06	0.06	0.14	0.22
12:00	0.1296	0.25	0.16	0.08	0.08	0.18	0.28
18:00	0.1396	0.28	0.19	0.11	0.11	0.23	0.34
24:00	0.1444	0.27	0.20	0.12	0.12	0.26	0.37
48:00	0.1499	0.26	0.20	0.12	0.12	0.27	0.38

Table 9. Standard deviation and error distribution for regression using the current wind speed, the average wind for the time which is being predicted, and the change in wind speed between the two times.

only a marginal improvement in accuracy as seen in table 9. Table 10 gives the coefficients from this regression. All the coefficients are highly significant with t-test values greater the 7. Again the constant (TConst) for the time for which the prediction is made becomes more significant for the longer prediction times.

A Fourier analysis of the wind speed data has been undertaken. As the data contained missing values this was done using regression of the individual cos and sin components. This also allowed an exact period of 24 hours to be used. The resulting power spectrum is plotted at the top of figure 6. Two clear peaks are seen and correspond to a 24 hour period (the left peak) and a 12 hour period. There is a possible indication of a small component with an 8 hr period. The bottom part of figure 2 is an artificial sequence generated from:

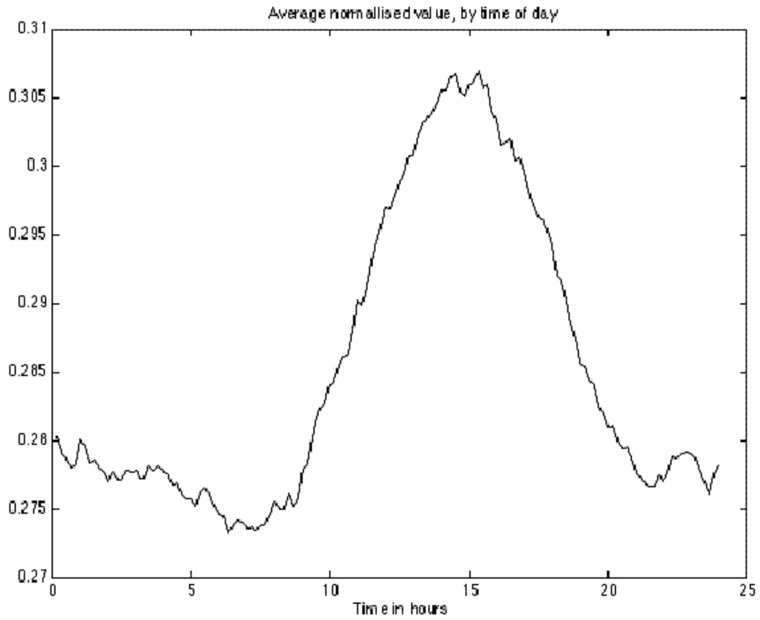

Figure 5. Average daily normalised wind speed for wind farm data.

$$p(i+1) = 0.99p(i) + N(0,1)$$

where $N(0,1)$ is a Gaussian random sample. 2^{19} elements of this sequence were generated and the power spectrum obtained by fast Fourier transform. It can be seen that the random sequence is essentially the same as the wind data except for the two peaks. It can be noted that forward predictions for this sequence are made by taking a multiple of the current value which reduces to zero as the time difference increases, and for large time intervals the prediction goes to the mean value for the sequence (which is zero in this case).

To examine the daily behaviour more closely clustering was used. To reduce the number of variables in a day the data was converted to an average hourly values and only data that was complete for the day was used. This provided 1154 days of complete data. A simple K-means clustering was used with the number of clusters being specified. The distance measure used was calculated by first determining the factor that brought the wind speeds for one day into closest agreement with those of the cluster mean (by least squares estimation of the factor), and

Hr:Min	Coef	TConst	Time diff
0:10	0.9886	0.0114	0.4193
0:20	0.9777	0.0224	0.7599
0:30	0.9681	0.0321	0.8456
1:00	0.9414	0.0588	0.9113
2:00	0.8926	0.1077	0.8896
3:00	0.8469	0.1536	0.8421
6:00	0.7212	0.2792	0.7151
12:00	0.5233	0.4771	0.5323
18:00	0.3965	0.6037	0.3997
24:00	0.3117	0.6878	0.0000
48:00	0.1690	0.8317	0.0000

Table 10. Coefficients for regression using the current wind speed, the average wind for the time which is being predicted, and the change in wind speed between the two times.

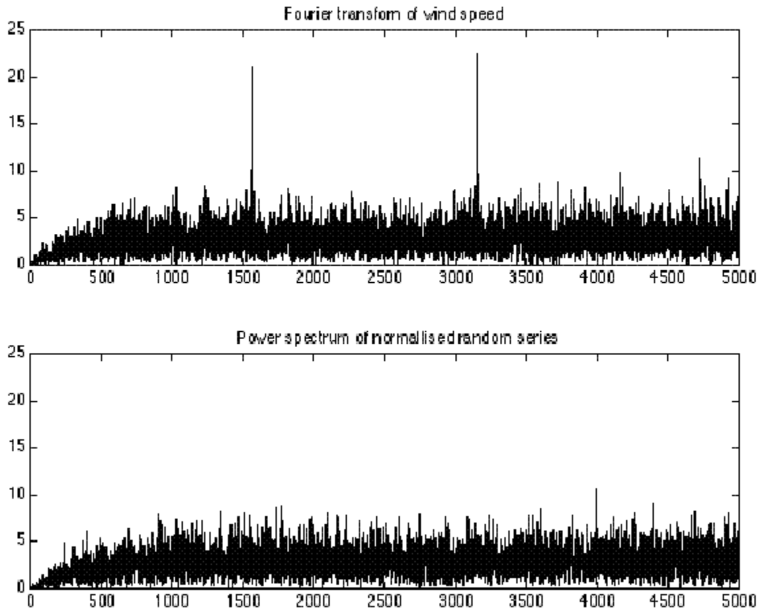

Figure 6. Power spectrum for wind data (top) and for a randomly generated sequence (bottom).

then determining the root mean square of the differences in the hourly values. For these clusters no attempt was made to identify outlying data that might allow better definition of the central part of the cluster.

Cluster	S.D.	Proportion
1	0.0476	0.0589
2	0.0606	0.0936
3	0.0671	0.3674
4	0.0644	0.4801

Table 11. Standard deviation of error in wind values when it is known which daily cluster and scaling factor should be used. Also the fraction of days allocated to a given cluster.

Different numbers of clusters were tried, up to seven clusters could be identified however four main wind profiles clusters were thought to be a reasonable choice. Figure 7 gives the means of the four clusters. The standard deviations and fraction of days allocated to a given profile are given in table 11. The standard deviations are significantly lower than the data mean of 0.1525 and better than the regression predictions for two or more hours ahead. Figure 8 shows the fit of the data after scaling to the cluster means. These graphs still contain a significant amount of unexplained or random variation.

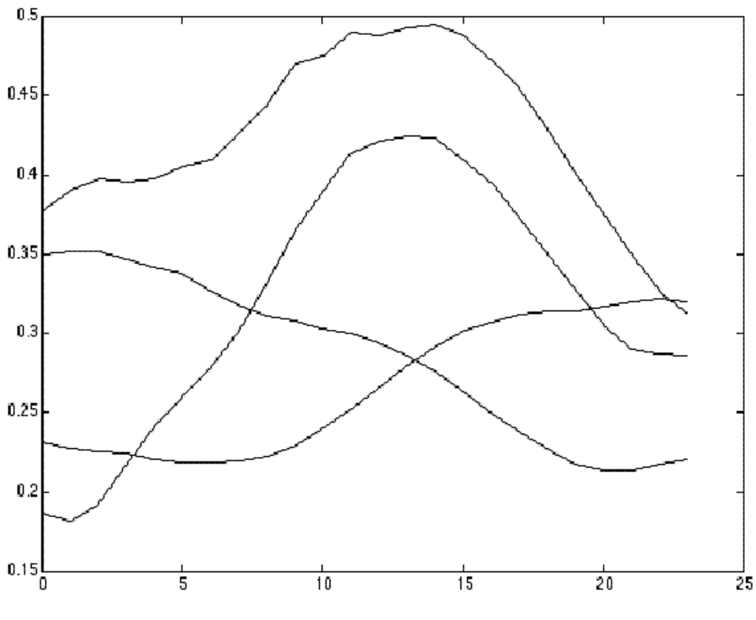

Figure 7. mean values for four clusters that give different daily wind profiles.

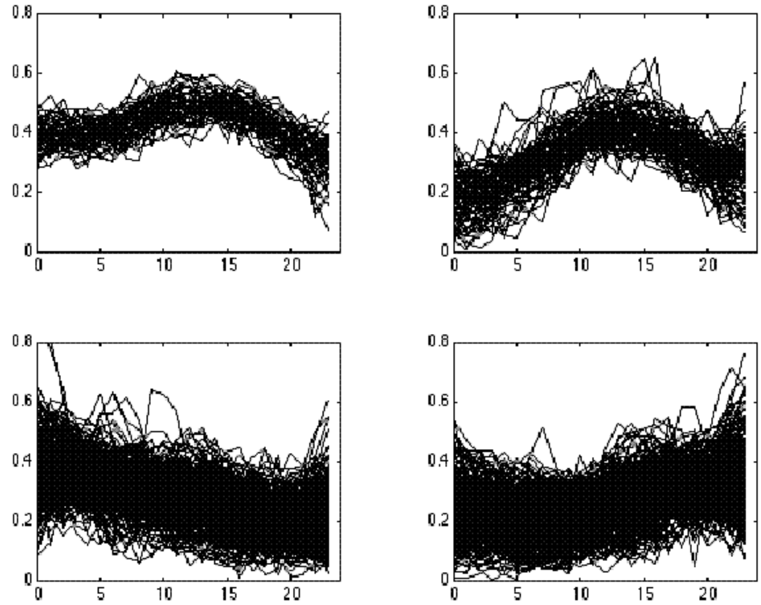

Figure 8. Fit of daily wind profiles to the mean profile of the four clusters.

The prediction of wind speed could be considerably improved if it can be determined which daily profile should be used. It is not known to what extent meteorological information can provide information on the best profile to use. Only limited meteorological data is included in the current data set. Figure 9 shows the average daily wind vectors (ends of the vectors giving speed and direction are plotted). Clusters one and two are do not occur at low wind speeds, but otherwise the clusters overlap on these diagrams.

The sequence of cluster occurrence has been investigated by calculating the probability of the transfer to a cluster on the following day, given the current cluster. Table 12 gives these probabilities.

Taking into account the number of days (890) used in calculating this table this is not far from the variation expected due to random variation from the proportions for all the clusters combined (Chi-squared p-value of approximately 0.84). It is thus considered that knowing the wind profile from the previous day provides little information on the profile for the current day. There may be advantage in using profiles for predictions of less than a day ahead, however this has yet to be investigated.

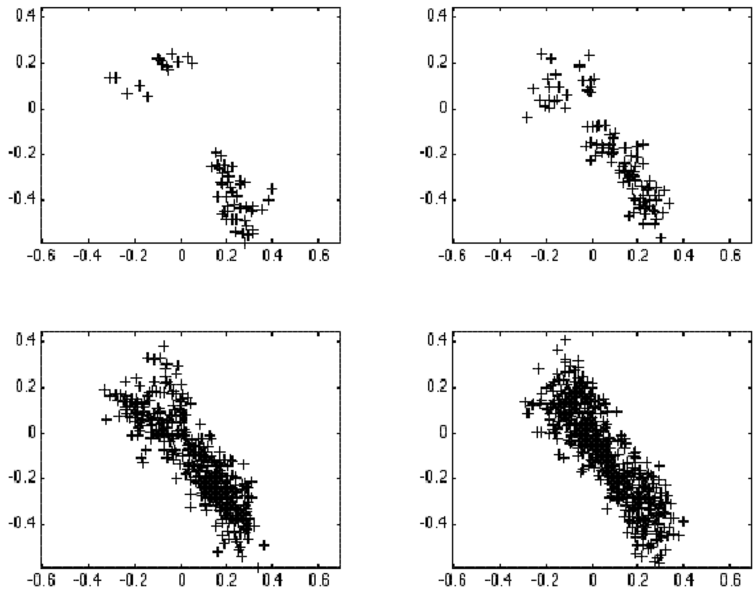

Figure 9. Wind velocity vectors (origin to points marked +) for the four clusters.

Comments on Regression Analysis of Wind Farm Data

A simple regression (current wind and a constant) is better than the persistence method, particularly as the time interval increases. This is due to the regression being able to use a constant in the regression, and thus transfer the prediction from relying on the current wind speed to relying of a constant at longer time intervals.

Regressions using the daily average profile and the current wind speed, as a correction gave only very minor improvements to the simple regression prediction.

Fourier analysis demonstrated a daily variation, and except for this the remainder of the spectrum seems indicate random variation only.

Wind direction could be added to the regression, or the components of the wind velocity used. This has not been tested in detail but is not expected to have a major effect on accuracy.

It is not known how much combining meteorological wind speed forecasts with the regression will improve the forecast for the wind farm. This is possibly the best regression option available to obtain better predictions. Good results were obtained using the meteorological pre-

Current cluster	Probability 1	in cluster 2	next 3	day 4
1	0.0449	0.2028	0.4038	0.3486
2	0.1063	0.1007	0.3117	0.4812
3	0.0393	0.0970	0.3492	0.5144
4	0.1105	0.0976	0.3707	0.4212
All	0.0589	0.0936	0.3674	0.4801

Table 12. Given the current day wind profile is in the cluster given by the rows the columns give the probability of the wind profile on the next day being in the cluster given by the columns.

	ANN	Persistence Method	Improvement by ANN
4-hour forecast Mean Absolute Error	0.1738	0.2007	0.0269 (13%)
12-hour forecast Mean Absolute Error	0.2592	0.3097	0.0505 (16%)

Table 13. Comparison of average forecast error.

dictions in regression as described in the section on the meteorological data.

The data can be divided into four (or more) typical daily profiles. Developing this further provides one promising approach. It is likely that the full meteorological forecast information will give sufficient information to reasonably predict the wind speed profile for several hours ahead, or the profile might be identified from previous wind speed and direction values. Combining this with the local speed data should give improved prediction.

5. The Artificial Neural Network (ANN) approach to wind power forecasting

An Artificial Neural Network (ANN) (e.g. Herve 1999, Picton, 2000) is a technique that is used to map any random input to a corresponding random output without assuming any fixed relationship between them. Neural Networks utilize past data in order to recognize a hidden pattern so that it can be used to forecast future values. A ANN study with a Multilayer Perceptron model was carried out to forecast wind power 4 hours and 12 hours ahead (figures 10 &11).

To compare the performance of the two methods in terms of the average forecasting error, all errors were summed up weighted by their frequency of occurrence. The comparison of average forecast error is shown in table 13.

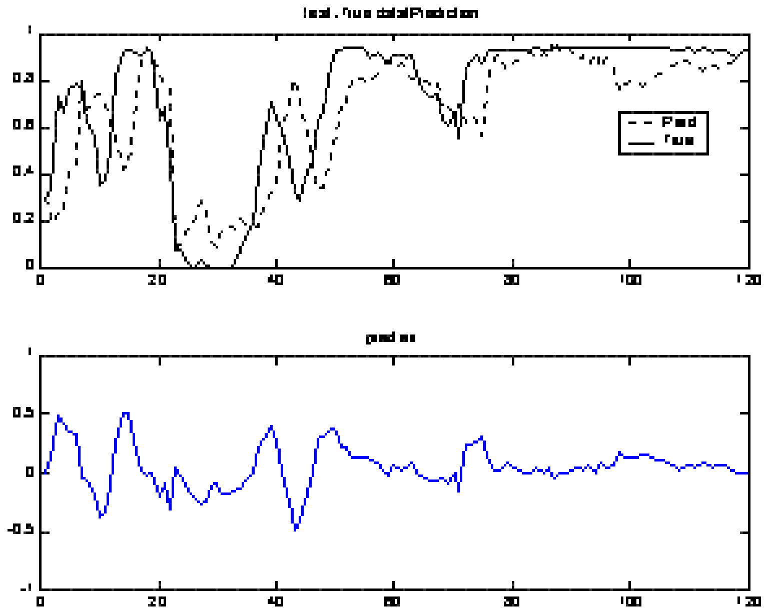

Figure 10. A sample 4-hour ahead ANN forecast.

Table 13 shows that in 4-hour and 12-hour forecast, ANN reduces the average forecast error by 13% and 16% in 4-hour and 12-hour forecasts respectively.

It was suggested that the adoption of two different models, one for daytime and one for night time could significantly improve the forecasts.

A similar ANN study utilizing a 3-layer perceptron was performed to examine the forecast of wind speed 4 hours ahead. On average, the predicted wind speed 4 hours ahead was about right, but there is a wide spread around that average. 95% of actual values are within 0.2 of the forecast. An evaluation was done using data not used in the modelling. Figure 12 is the histogram of the residuals and figure 13 is the cumulative frequency of errors.

6. The Kalman Filter

The Kalman Filter (Anderson 1979, Jazinski 1970) is useful in situations where at some time t the value of a variable may be predicted k

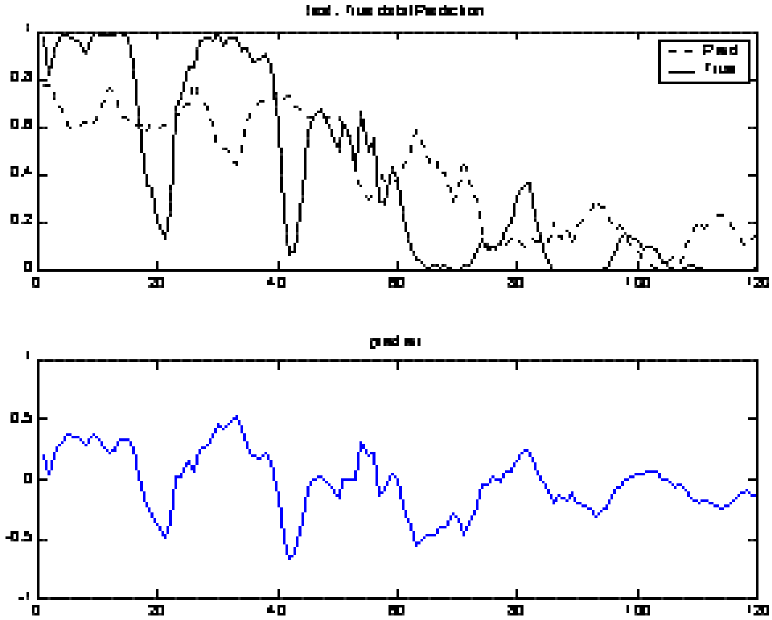

Figure 11. sample 12-hour ahead ANN forecast.

steps ahead, given its last N recorded values. If ν_{t+k} is at time t, the predicted speed at time $t+k$, it can be written as linear combination of the last N measured speed values:

$$\nu_{t+k} = a_t k_t + a_{t-1} k_{t-1} + ... + a_{t-N+1} + \nu_{t-N+1}$$

where the coefficients, $a_t, a_{t-1}, ..., a_{t-N+1}$, are generally variable in time, and N is the order of the filter. The forecast may be one or more steps ahead. The Kalman Filter is an algorithm whereby the coefficients, a_{t-i} , are calculated from the previous a_{t-i-1} , $i = 0, 1, ...N-1$, using a relationship which depends on the latest measurements of ν.

As a demonstration a simple Kalman filter with only one term in the prediction equation was tested. The results are given in table 14 and figure 14. It can be seen that a simple Kalman filter is capable of providing reasonable predictions. It is expected that one or two additional terms would make a marginal improvement to the prediction quality, but time did not allow testing.

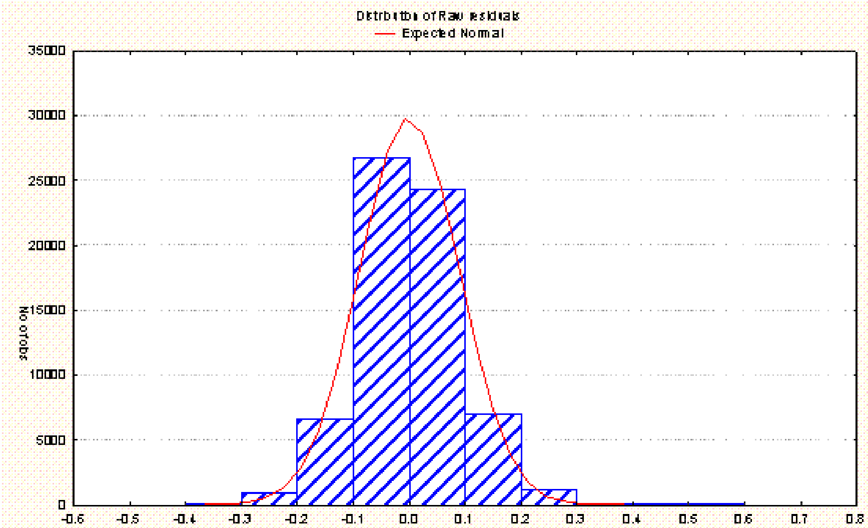

Figure 12. Histogram of the residuals.

7. Conclusions

The prediction of wind speed is known to be a difficult problem with most methods giving only a limited improvement over the simple persistence method.

The prediction of wind speed has the reputation of being the most difficult of the meteorological variables to predict. The analysis above indicates that there is a significant apparently random component to the wind speed. Fourier analysis found only a daily systematic component in the frequency spectrum.

There is a daily average pattern of wind speed, however its use provides only a marginal improvement in predictions for the wind farm data. Several different daily patterns can be found within the speed data and these, if it is known which applies, can assist in forecasting. It is not clear how accurately meteorological conditions can determine which pattern applies at a given time.

Regression techniques provide a clear description of the nature of the prediction, but are dependent on having appropriate variables available to make the prediction from. They have the advantage of being well understood and being backed by a considerable amount of statistical knowledge.

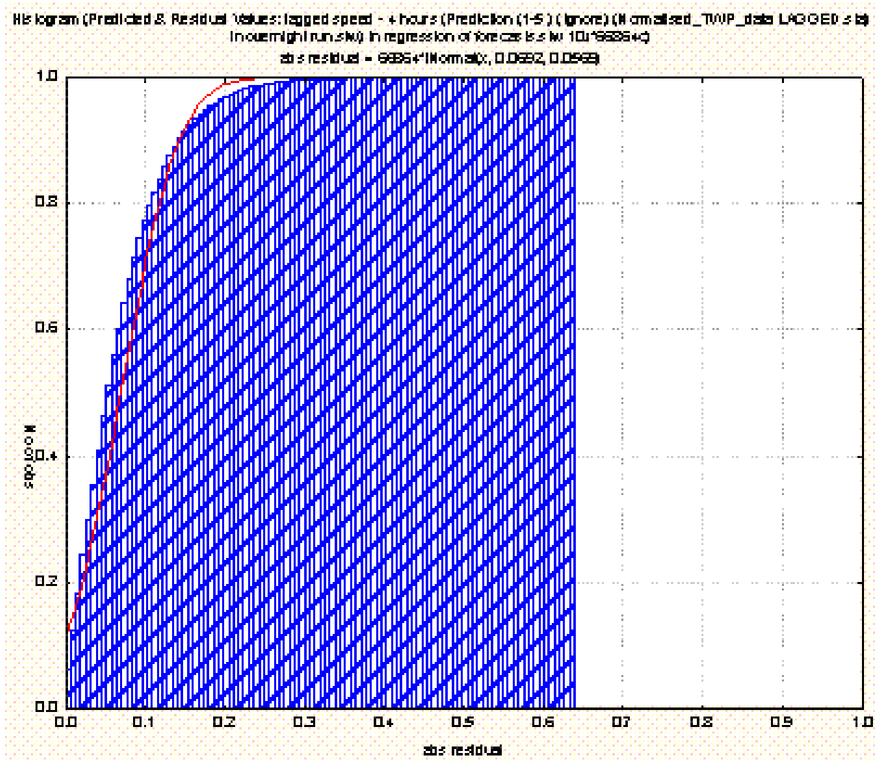

Figure 13. Cumulative frequency of errors.

Neural networks provide a black box method of using many possible inputs to the prediction and an automatic creation of a prediction. However the logic behind the prediction is not available, and there is the possibility that in some circumstances the prediction may be quite unreasonable. It is very difficult to check all aspects of the behaviour of a neural network with many inputs. The results obtained are not directly comparable with the regression results, but do show that neural networks can produce competitive predictions.

The Kalman filter is basically a regression technique placed on top of a dynamic model, as opposed to a static model in simple regression. The advantages of the Kalman filter rely heavily on the quality of the dynamic model. It is not clear how appropriately the simple linear models often used with the Kalman filter apply to the forecasting of wind

Wind Speed	Prediction	Difference
0.4198	0.3092	0.1107
0.4018	0.4213	-0.0195
0.3976	0.4016	-0.0040
0.3258	0.3975	-0.0717
0.2904	0.3249	-0.0345
0.3353	0.2899	0.0454
0.3429	0.3359	0.0070
0.3310	0.3430	-0.0120
0.3447	0.3308	0.0139
0.3447	0.3449	-1.81E-04

Table 14. Sample output from a Kalman filter.

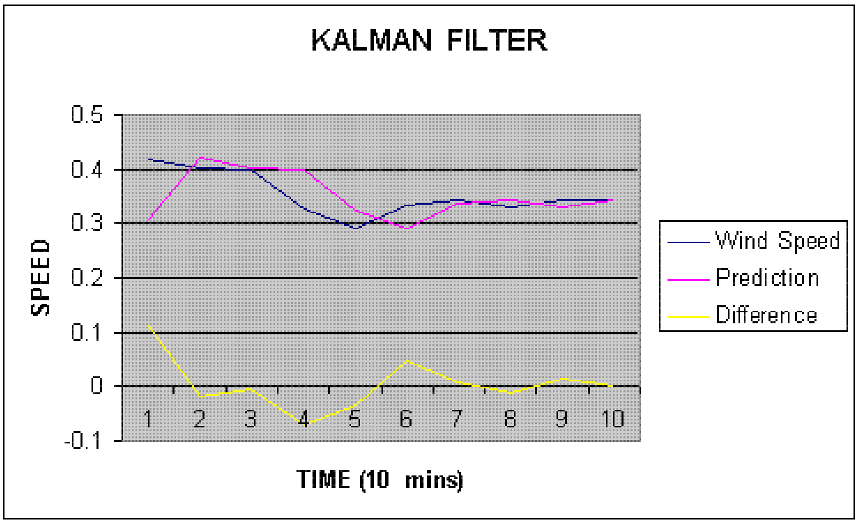

Figure 14. . Kalman Filter tracking of actual wind speed. 1 step ahead (1 step = 10 minutes).

speed. A forecasting model with input of meteorological conditions may be a more useful approach.

The meteorological data made available was disappointing in the accuracy of its wind speed predictions. Unfortunately the values recorded were instantaneous wind speed, which undoubtedly added some additional random variation. It seems that local wind speed over short periods is not closely related to the larger scale of the forecast data. It is believed that wind direction and rainfall are more reliably predicted in forecasts than wind speed where turbulence apparently adds a significant random component.

A combination of meteorological predictions and regression gave better predictions than either value alone when tested on the meteorological data. It was not possible to combine the meteorological data with the wind farm data as they were taken at different times (as well as different locations and different averaging).

The best approach to predicting the wind speed would seem to come from using a combination of both local measurements and meteorological data. The local measurements can be at the wind farm site, but there would seem to be considerable potential benefit from measurements in the region of the wind farm. The meteorological data could be used in several forms. There is the direct prediction of wind velocity from the numerical forecast, information from the pressure and wind direction forecasts (isobaric charts) over the region, and/or a matching of current conditions with those in a record of past weather. These alternatives need to be investigated to determine which gives the best results.

Acknowledgements

The project coordinators wish to thank, the industry sponsors Conrad Edwards, Greg Williams, and Brian Kirtlan for their support and enthusiasm, and the project participants Kaye Marion, Barry McDonald (data preparation & initial regression), Ray Hoare, Zeke Chan, Timothy Hong (neural networks), Boda Kang (Kalman filter), Bruce Craven, John Cogill, Manju Agrawal, Ian Wright, and Andy Philpott. The work of the organisers of the MISG and particularly of Graeme Wake is also greatly appreciated.

References

Anderson, B.D.O., and Moore, J.B., 1979. Optimal Filtering, Prentice-Hall : New Jersey.

Giebel G. 2003. The state of the art in short term prediction of wind power, Project ANEMOS (http://anemos.cma.fr/), http://anemos.cma.fr/modules.php?name=Downloads&d_op=viewdownl

Herve, A., Valentin, D., and Edelman, B. 199. Neural Networks, Sage Puplications : California.

Jazwinski, A.H., 1970. Stochastic Processes and Filtering Theory, Academic Press : New York.

Picton, P., 2000. Neural Networks, Macmillan Press : New York.