A statistical analysis of the reduction of the wind power prediction error by spatial smoothing effects

Ulrich Focken, Matthias Lange [1], Kai Mönnich, Hans-Peter Waldl

Department of Energy and Semiconductor Research, Faculty of Physics
Carl von Ossietzky University of Oldenburg, D-26111 Oldenburg, Germany

Hans Georg Beyer, Armin Luig

Department of Electrical Engineering, University of Applied Sciences (FH) Magdeburg,
D-39114 Magdeburg, Germany

Abstract

We discuss the accuracy of the prediction of the aggregated power output of wind farms distributed over given regions. Our forecasting procedure provides the expected power output for a time horizon up to 48 hours ahead. It is based on the large scale wind field prediction which is generated operationally by the German weather service. Our investigation focuses on the statistical analysis of the power prediction error of an ensemble of wind farms compared to single sites. Due to spatial smoothing effects the relative prediction error decreases considerably. Using measurements of the power output of 30 wind farms in Germany we find that this reduction depends on the size of the region. To generalize these findings an analytical model based on the spatial correlation function of the prediction error is derived to describe the statistical characteristics of arbitrary configurations of wind farms. This analysis shows that the magnitude of the error reduction only weakly depends on the number of sites and is mainly determined by the size of the region. Towards a correction of systematic prediction errors an analysis of the temporal structure of the forecast error is performed. For this purpose the correlation of the errors for consecutive forecasts is analyzed for single sites and ensembles.

Key words: Wind power; Short-term prediction; Spatial correlation; Smoothing effects

[1] Corresponding author. Tel ++49-441-798-3577, Fax ++49-441-798-3326
Email: m.lange@mail.uni-oldenburg.de

1 Introduction

The contribution of wind energy to the public electricity supply has grown rapidly in countries like Denmark, Germany and Spain over the last decade. In some regional grids the installed capacity of wind turbines is of the order of magnitude of the minimal electrical load (approx. 30 % of max. load). In contrast to conventional electricity production the availability of wind energy depends on fluctuating meteorological conditions posing new challenges to grid operators and utilities. The feed in of electricity from wind turbines acts as a negative load leading to increased variations of net load patterns. Due to the uncertainty of the temporal development of the wind fields additional spinning reserves have to be kept ready in conventional power plants to replace the wind energy share in case of suddenly decreasing wind speeds. To permit an efficient integration of the fluctuating wind energy into the existing supply system the expected wind turbines' output has to be known in advance. For this purpose different wind power prediction systems with a time horizon of up to 48 hours have been established in recent years.

A common approach to short term wind power prediction is refining the output of numerical weather prediction (NWP) models operated by weather services to obtain the local wind conditions. This concept was used to develop reliable operational prediction systems in Denmark [1] and Germany [2,3]. Prediction tools using advanced statistical methods are operated on a daily basis at a utility in Denmark [4] (recently been replaced by a new system that also includes input from NWP models [5]) and are under development in Ireland [6].

A major issue regarding the operational use of wind power prediction tools is, of course, the quality of the forecast. As the accuracy of a given NWP model output is not known the uncertainty of the wind field is directly transfered to the power prediction procedure. Unfortunately, the nonlinearity of the wind turbines' power curves leads to a further amplification of the error, i.e. small deviations in the wind speed may result in large deviations in the power. For single sites this can lead to considerable prediction errors. However, in operational use a wind power prediction system is supposed to provide the combined power production of many wind farms installed in a region with some 100 km in diameter. The error of the prediction of this regional power output is expected to decrease due to spatial smoothing [7]. In this paper we focus on the two major variables determining the magnitude of this statistical effect which are the spatial extension of the region and the number of wind farms it contains.

For comparison, we use measurements from 30 wind farms in Germany to form typical regions with different extensions corresponding to a medium and large utility supply area and sum up the according measured power output. Fictitious model ensembles together with the correlation function based on the measured data allows us to shed some light on the general statistical behaviour of distributed wind farms

2

regarding the prediction error. As a special case the reduction of the prediction error is calculated for the real distribution of all German wind farms.

2 Forecasting method

The principle scheme of our prediction system can be seen in figure 1, for a detailed description see [2,3]. As input the result of an operational numerical weather prediction model is used. The German weather service (DWD) currently operates the "Lokalmodell" which replaced the "Deutschlandmodell" in November 1999. Two main runs are started every day at 00 UTC and 12 UTC. Our calculations are based on the wind speed and direction forecast up to 48 hours coming from the "Deutschlandmodell" (00 UTC run). The resolution of the data is 14×14 km^2 (newer NWP models use resolutions up to 5×5 km^2), i.e. rather sparse, so a spatial refinement is necessary to predict the wind power at a specific site. We calculate the wind speed at hub height using the geostrophic drag law and the logarithmic wind profile. Thermal stratification of the boundary layer, roughness around the site, orography and wind farm shadowing effects are included.
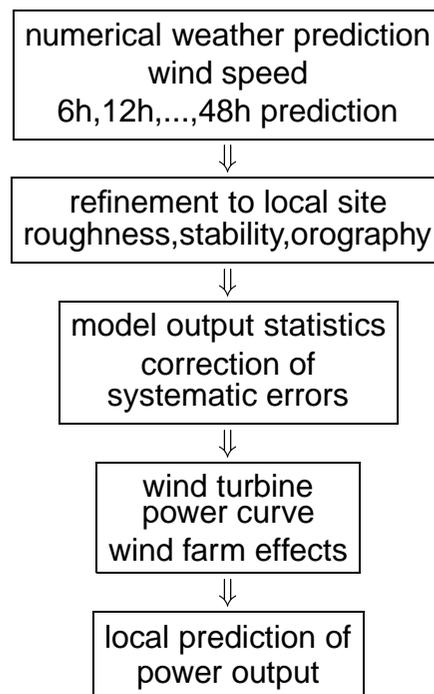

Fig. 1. Principle of the spatial refinement of the numerical weather prediction leading to a local prediction of wind conditions and wind power.

3 Prediction error of single sites

The accuracy of the power prediction for a single site is determined by comparing the results of the locally refined prediction with measurements [3].

For this purpose archived NWP data for the years 1996 to 1999 was provided by the German weather service. In particular, we use the 6, 12, 18, 24, 36, and 48 h predictions of the 00 UTC run. Measured data for the same period of time was collected in the framework of the German Scientific Measuring and Evaluation Programme (WMEP) carried out by ISET, Kassel [8].

Figure 2 shows a comparison between prediction and measurement for the power output of a wind turbine in the North German coastal region. In general the predicted and the measured time series show a good agreement. Significant differences can be seen mainly for the 36 and 48 h prediction. In particular, the beginning of a storm on day 326 is not correctly predicted and on day 330 the prediction shows a time shift of several hours.
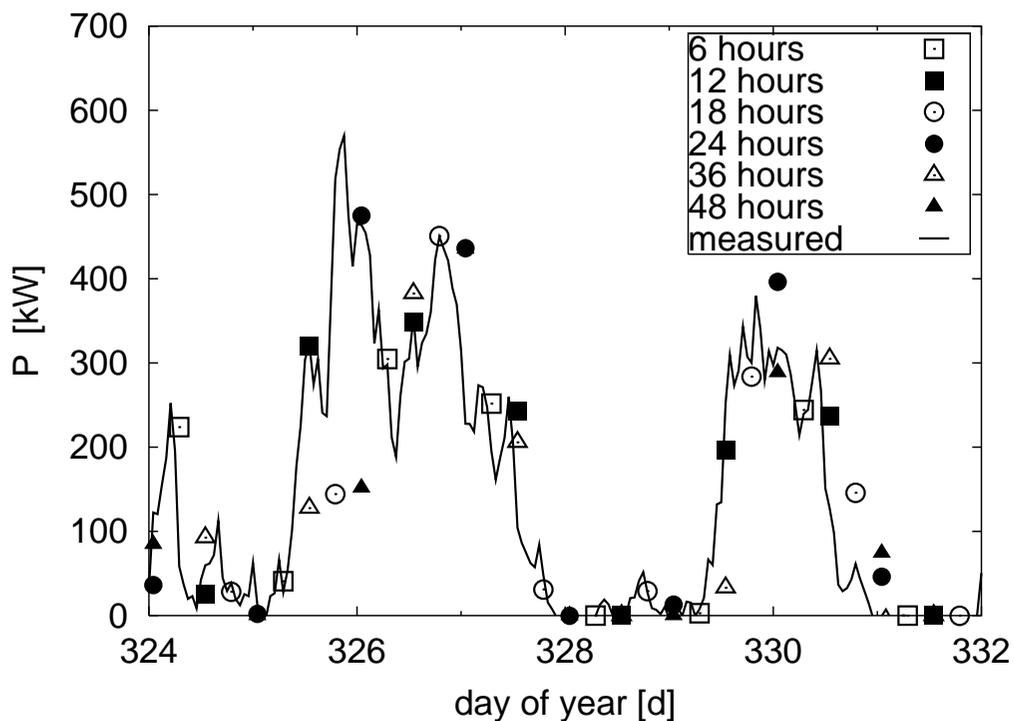
Fig. 2. Typical time series of measured and predicted power output for one site. There is a good agreement between prediction and measurement. On day 325 the beginning of a storm was not correctly predicted by the 36 and 48 h forecasts. The prediction of the storm around day 330 is several hours too late.

For our investigation the quality of the prediction is quantified by the standard deviation of the difference between predicted and measured power normalized to

the installed power P_{inst} of the wind turbines in the period of time to be considered:

$$\sigma = \frac{1}{P_{\mathrm{inst}}} \sqrt{\frac{1}{M} \sum_{i=1}^{M}[(P_{\mathrm{pred},i} - P_{\mathrm{meas},i}) - (\overline{P_{\mathrm{pred}}} - \overline{P_{\mathrm{meas}}})]^2} \qquad (1)$$

P_{pred} is the predicted power output, P_{meas} the measurement and M the number of data points. We choose the standard deviation, σ, instead of the commonly used root mean square error (rmse) in order to compare the results to a statistical description based on a correlation function (section 4.2) which naturally involves σ. The difference between these two error measures is that the bias, i.e. $\overline{P_{\mathrm{pred}}} - \overline{P_{\mathrm{meas}}}$, is not subtracted in the rmse. In our case this bias is quite small leaving the rmse values only slightly higher than the σ values.

Figure 3 summarizes the results of the comparison between measured data and predictions for single sites. The standard deviation normalized to the installed power and averaged over 30 wind farms rises from 12% for the 6 h prediction to 18% for 48 h. The increase of the prediction error with increasing time horizon might be due to the growing systematic error in the numerical weather forecast for longer prediction times. For comparison figure 3 also shows the performance of the persistence forecast, i.e. the prediction for a point of time in the future is set to the current measurement. For short prediction times below 6 h the persistence is comparable to or even better than the prediction based on NWP models. With increasing prediction horizon persistence gives a significant higher error.

The correlation between predicted and measured time series measures the similarity of the two signals. As shown in figure 4 the correlation for our system decreases from 0.84 for the 6 h to 0.70 for the 48 h prediction while the persistence shows a rapidly decaying correlation.

4 Spatial smoothing

Under operational conditions a prediction of the combined power output of many wind farms distributed over a large region, e.g. the supply area of a utility, is needed. By integrating over a region the errors and fluctuations underlying the measurement and the forecast at single sites cancel out partly. These statistical smoothing effects lead to a reduced prediction error for a region compared to a local forecast. The size of the region and the number of sites it contains are the main parameters that determine the magnitude of the error reduction. The analysis of measured data very clearly shows this effect but is constrained to a fixed ensemble of sites. To generalize our findings we use model ensembles which require a statistical description of the regional prediction error in terms of spatial correlations.
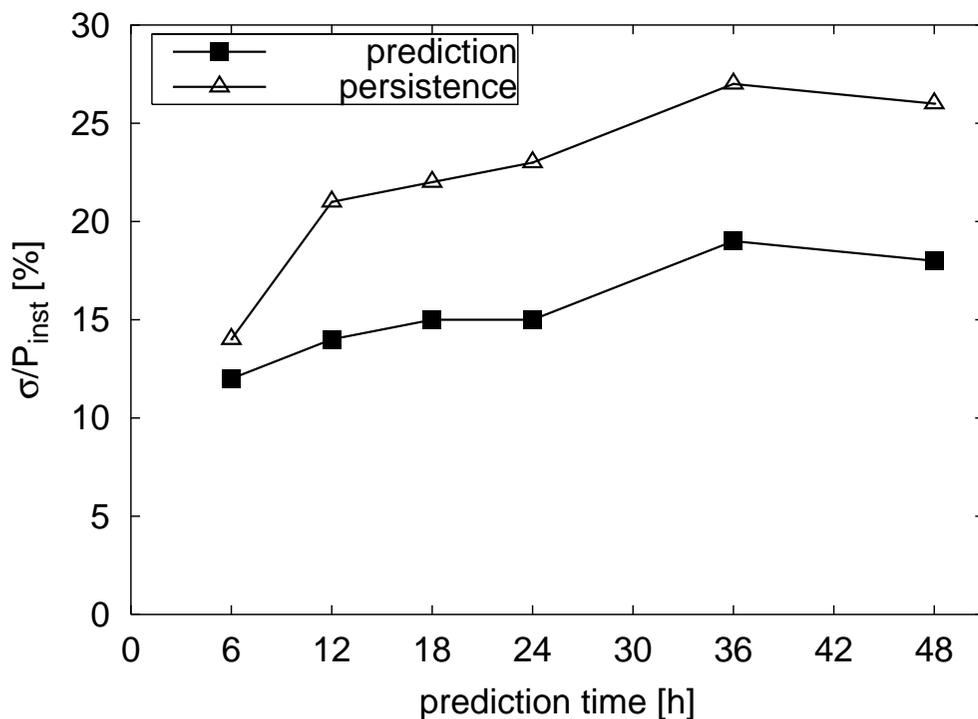
Fig. 3. Normalized standard deviation versus prediction time averaged over 30 wind farms for our prediction system and persistence. For prediction times below 6 h the error of the persistence forecast is smaller than or comparable to the prediction. For larger times the prediction system performs much better.

4.1 Ensembles of Measurement Sites

Our first approach is to investigate the spatial smoothing effect using data from an ensemble of 30 wind farms in the Northern part of Germany. The sites are divided into regions of two different types according to typical areas covered by a medium and a large utility. The smaller regions with a diameter of approximately 140 km (see figure 5) contain three to five measurement sites each. The larger regions are about 350 km in diameter with five to seven sites each. For comparison we form a very large region containing all sites which has a size of about 730 km.

The predicted and measured power output of a region is calculated by summing up the time series for every wind farm located in the region and dividing them by the number of wind farms. Similar to the analysis of single sites the standard deviation of the difference between these two ensemble time series gives the regional prediction error. Figure 6 shows the results for the different region sizes and various prediction times. The standard deviation, $\sigma_{\mathrm{ensemble}}$, of the ensemble, i.e. the regional prediction error, is normalized to the mean standard deviation of the single sites, σ_{single} according to section 3, and averaged over all regions of the same size. For the given ensemble this ratio decreases with increasing region size, e.g. the 6 h prediction gives an average ratio of 0.79 for the 140 km region, 0.68 for the 350 km
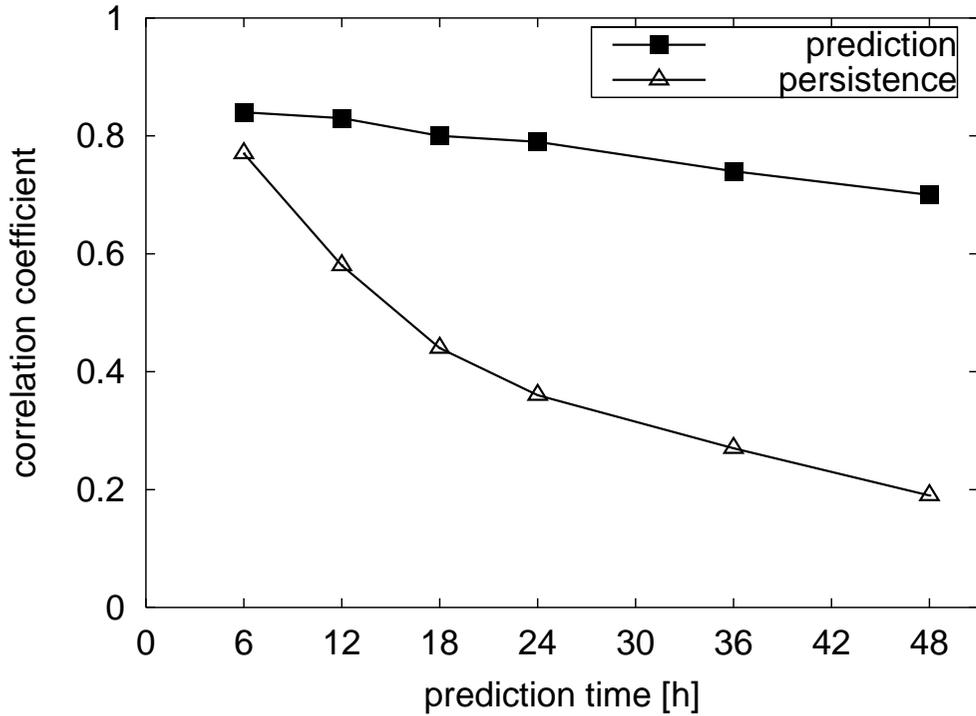
Fig. 4. Correlation of predicted and measured power time series over prediction time for our prediction system compared to persistence. Again persistence performs well for small prediction times but its correlation decreases very rapidly while for the prediction system the correlation remains on a rather high level.

region, and 0.50 for the 730 km region (figure 6). In all cases the reduction of the regional prediction error is less pronounced for larger prediction times.

4.2 Model ensembles

The analysis for the specific set of measurement sites shows a significant decrease of the prediction error compared to a single site. In order to draw general conclusions concerning other configurations of wind farms we analyze random ensembles of sites. This allows us to vary the size of the regions and the number of wind farms over large ranges to see how the reduction of the error depends on these parameters. For this purpose we need a proper statistical description of the regional prediction error.

The key element connecting the spatial distribution of sites with the regional prediction error are the cross-correlation coefficients r_{xy} of the difference between prediction and measurement, i.e. $P_{\text{pred}}(t) - P_{\text{meas}}(t)$, for the single sites. If r_{xy} is known for all pairs of sites, the standard deviation, σ_{ensemble}, of the differences between measurement and prediction can easily be calculated using the σ of the
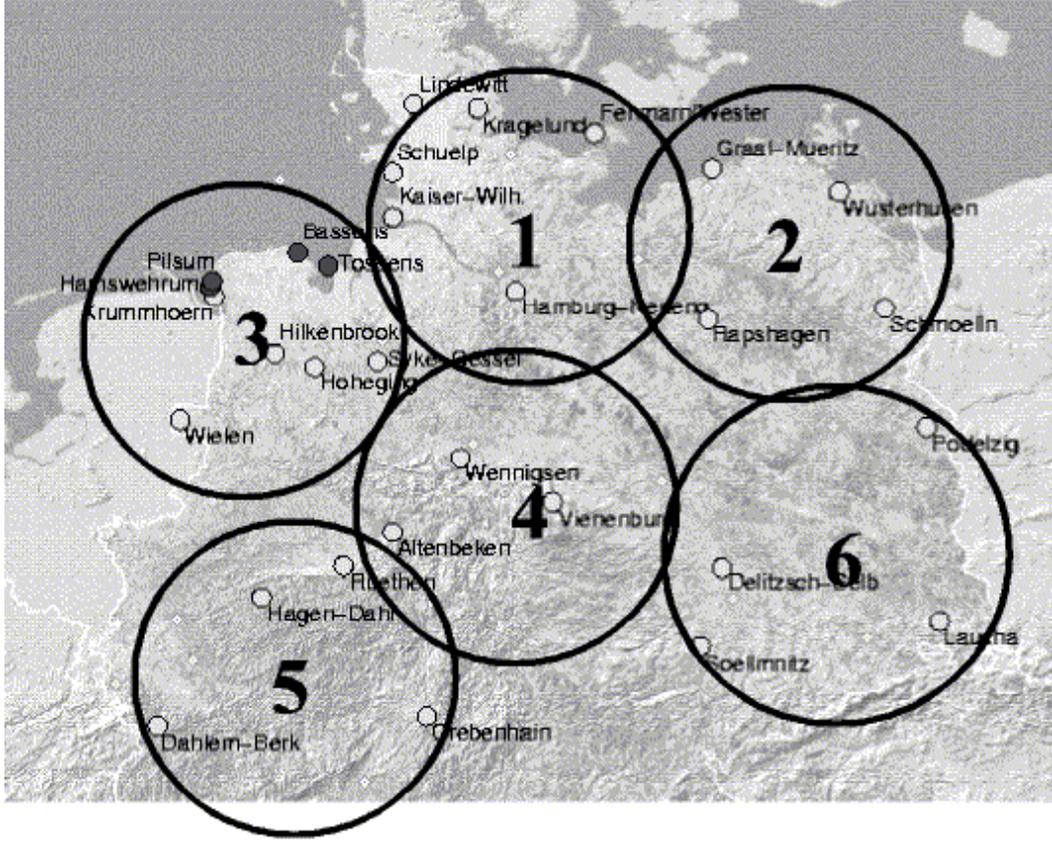
Fig. 5. Regions in Northern Germany with 140 km in diameter. The points denote the measurement sites.

individual sites by

$$\sigma_{\text{ensemble}}^2 = \frac{1}{N^2} \sum_x \sum_y \sigma_x \sigma_y r_{xy} \tag{2}$$

where N is the number of sites in the region, σ_x the standard deviation of the single sites.

To obtain an analytic function describing the dependence of r_{xy} on the site distance the following procedure is applied. For each pair of the 30 wind farms the cross-correlation coefficient between predicted and measured time series is calculated and ordered according to the distance between the two sites x and y. Figure 7 shows cross-correlation coefficients versus distance for the various forecast horizons where the data points have been averaged over 25 km bins. For small prediction times (6 and 12 h) the cross-correlation decreases rather rapidly within 150 km while for longer times (36 and 48 h) the decrease is much slower. This might be due to the growing systematic errors for increasing forecast horizon which give rise to higher spatial correlations.

Note that the cross-correlation coefficients of the measured power alone decay con-
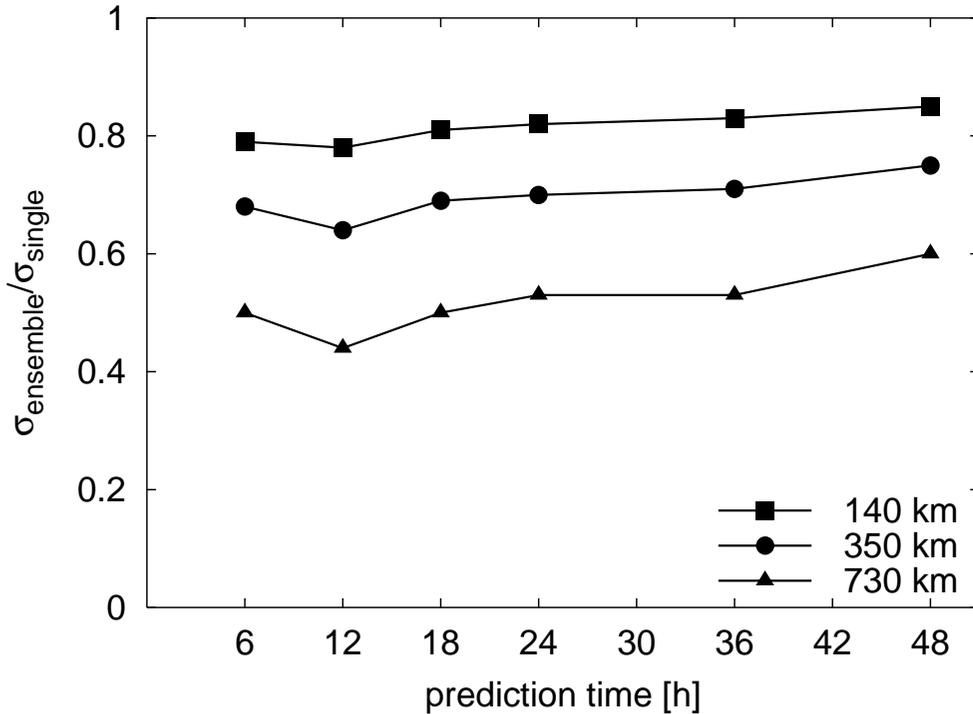
Fig. 6. Ratio between standard deviation of ensemble and single time series ($\sigma_{\mathrm{ensemble}}/\sigma_{\mathrm{single}}$) for various region sizes and forecast horizons for the ensemble of measurement sites. $\sigma_{\mathrm{ensemble}}/\sigma_{\mathrm{single}}$ decreases with increasing region size. In all cases the reduction of the regional prediction error is less pronounced for larger prediction times.

siderably slower with increasing distance than the cross-correlation of the deviations. As shown in [9] the cross-correlation of the power output of wind turbines is about 0.7 at 130 km while at the same distance r_{xy} drops to approx. 0.2 for the 6 h prediction (figure 7).

We obtain a proper description of the cross-correlation by fitting analytic functions of the form $r_{xy} = a \cdot e^{-d/b}$ (a and b are fit-parameters and d is the distance between the two sites) to the cross-correlation coefficients derived from the measured data. It turns out that piecewise exponentials lead to a suitable fit to the data points. With the correlation function r_{xy} based on the fitted data we can now use equation (2) to calculate the prediction error, $\sigma_{\mathrm{ensemble}}$, of the model regions. We normalize the σ_x of the fictitious wind farms to 1 which means that they all have the same weight.

The geographical coordinates of the model ensembles are chosen randomly. Each result given in the following represents an average value over ten realizations of ensembles with fixed size and number of sites.

Figure 8 shows the ratio between the regional error and the mean of single sites $\sigma_{\mathrm{ensemble}}/\sigma_{\mathrm{single}}$ for two regions with different sizes versus the number of sites in the region. The cross-correlation function r_{xy} based on the 36 h forecast was used. Obviously, $\sigma_{\mathrm{ensemble}}/\sigma_{\mathrm{single}}$ approaches a saturation level for increasing number of
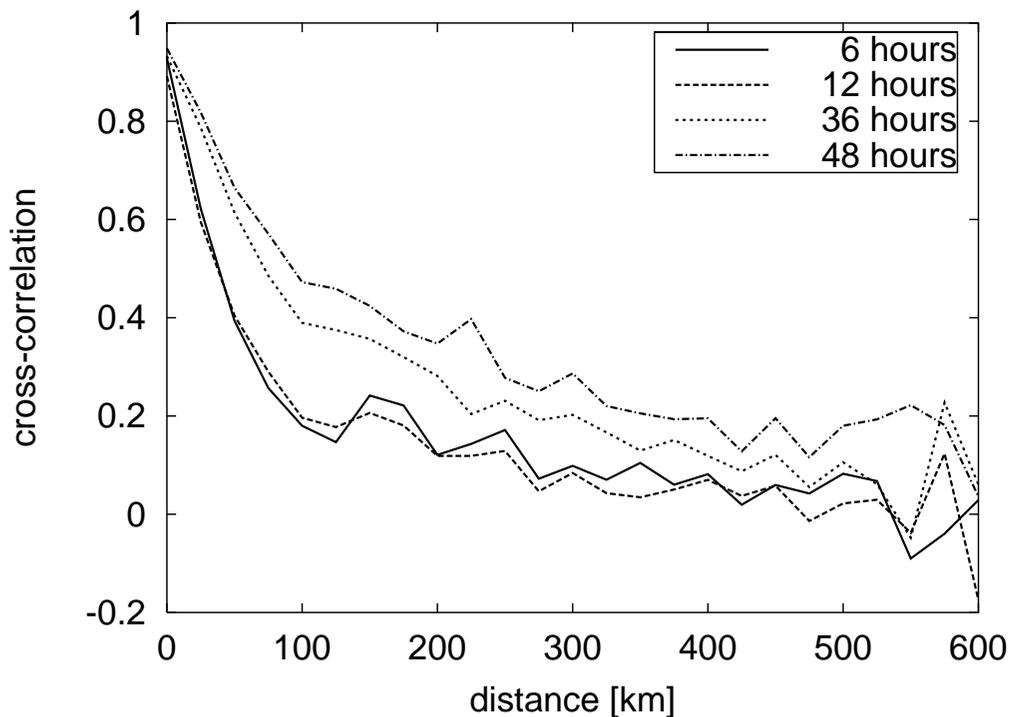
Fig. 7. Spatial cross-correlation of prediction deviations for various prediction times. The cross-correlation coefficients have been averaged over 25 km bins.

wind farms. This limit is already reached for a rather small number of wind farms. Beyond that point the error reduction does practically not depend on the number of sites, e.g. for the size of a typical large utility (approx. 370 km) less than 50 sites are sufficient to have a constant level of 0.63.

The experimental data of section 3 agrees very well with these calculations. Comparing the 36 h data in figure 6 with the values corresponding to the same number of sites in figure 8 we find $\sigma_{\mathrm{ensemble}}/\sigma_{\mathrm{single}}$=0.71 for the 360 km region for both ensembles. For the 730 km this ration is 0.53 for the ensemble of measurement sites and 0.52 for the model ensembles.

As expected the saturation level decreases with increasing size of the region. This is illustrated in figure 9 where the limit values for regions with different extensions containing 4000 sites are shown. There is a rapid decay for extensions below 500 km.

The distribution of distances within the ensemble of wind farms plays a crucial role in explaining the saturation level. Figure 10 shows typical frequency distributions of the distances in a 360 km region for different numbers of randomly distributed wind farms. The distributions correspond to specific ensembles, i.e. another realization has a different distribution. Figure 10 clearly indicates that the distribution functions converge for an increasing number of wind farms. Beyond a certain number of wind farms in the region the characteristic distribution of distances approaches a

10
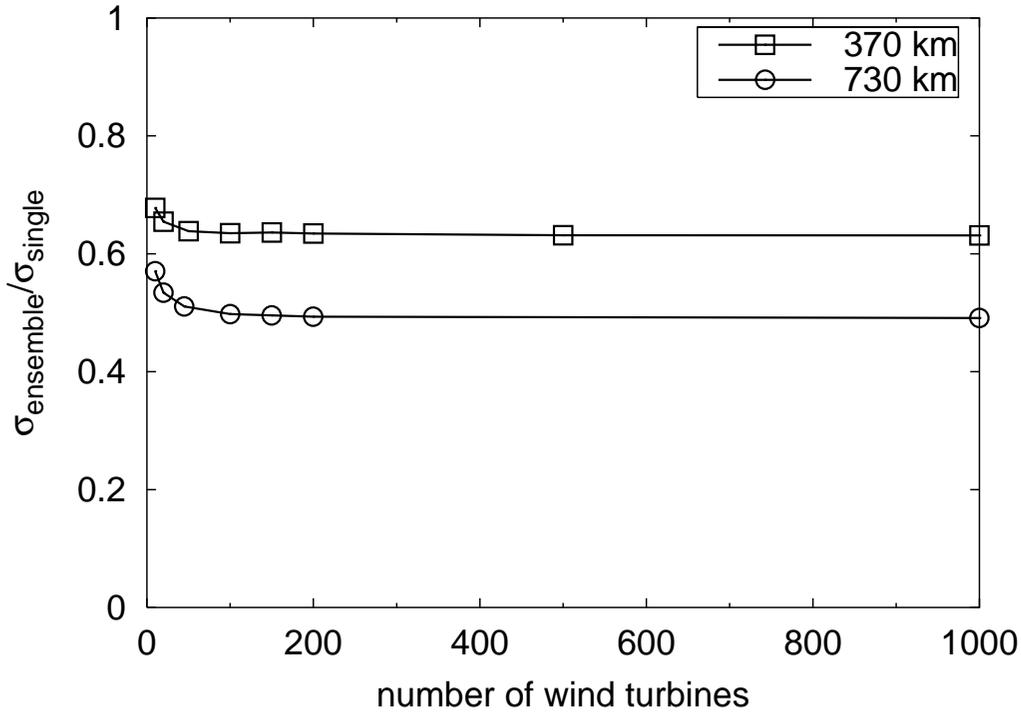
Fig. 8. Ratio $\sigma_{\text{ensemble}}/\sigma_{\text{single}}$ versus number of sites for the model ensembles. Each data-point represents an average over 10 ensembles. The fitted cross-correlation function for the 36h forecast was used. A saturation level is reached for a small number of sites.

limit distribution such that the regional error given by equation (2) does no longer change. If the region has a certain "population density" of wind farms it contains enough pairs to represent all possible correlations which contribute to σ_{ensemble}. Therefore, adding more sites to the region does not reduce the regional error, as would be expected for uncorrelated sites, and a saturation is reached.

The limit distribution has a mean value which is typically about half the diameter of the region. As the distribution is concentrated around the mean most contributions to σ_{ensemble} come from this interval. For increasing region size this leads to a decreasing saturation level due to the exponential decay of the cross-correlation function (figure 7).

4.3 Distribution of German wind farms

Finally, we consider the real distribution of the wind farms in Germany (in 1999) as a special model ensemble and compare the regional error to a single site as above. For the 36 h prediction this gives $\sigma_{\text{ensemble}}/\sigma_{\text{single}} = 0.43$. Note that this ratio for an equivalent region of the size of Germany with randomly distributed wind farms would be lower because the real ensemble shows a strong imbalance in the distribution of sites in the North and South (figure 11).

11
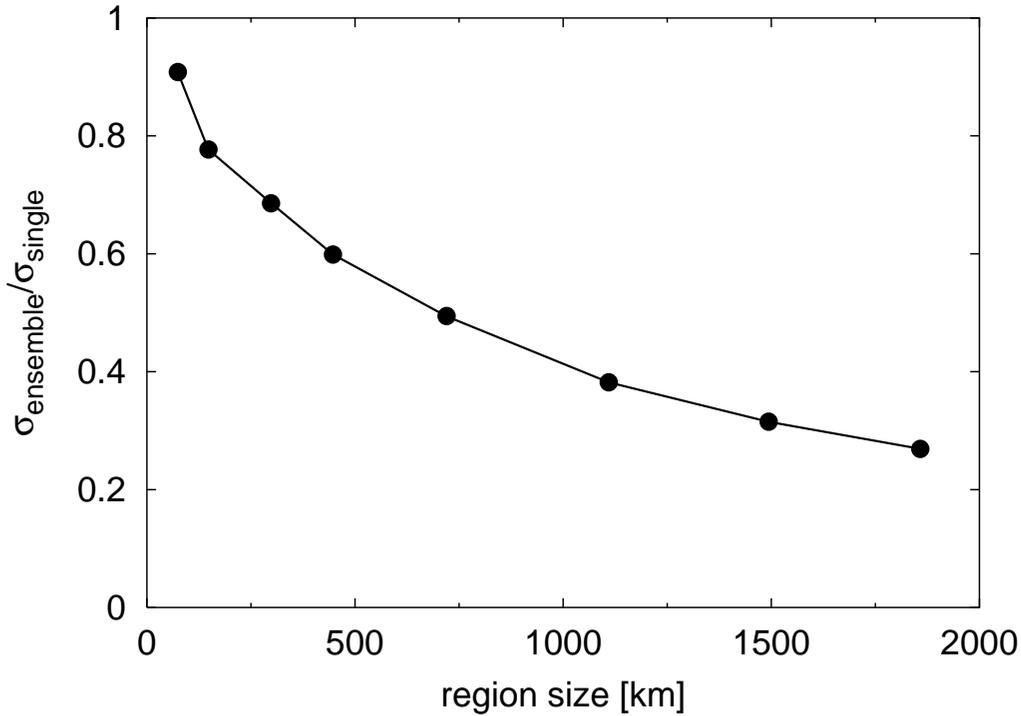
Fig. 9. Saturation values of $\sigma_{\text{ensemble}}/\sigma_{\text{single}}$ (4000 sites) for the 36 h forecast. Each data-point represents an average over 10 ensembles.

5 Analysis of the temporal structure of the forecast errors

For a further improvement of the forecast quality, methods to correct the actual forecasts using knowledge of comparison to previously measured data (so called model output statistics) are discussed (see e.g. [3]). A basic approach in this context is the analysis of the auto-correlation structure of the errors. A typical scatter diagram for the single site forecast errors of consecutive days is shown in figure 12 (left). The respective inspection shows that the auto-correlation coefficient is about 0.2 only. Thus the application of a simple linear correction procedure based on the deviation between prediction and measurement from the previous day will not lead to any remarkable improvements of the forecast.

Looking again for the change of the structure of the forecast errors when going from single sites to ensemble data gives the result that the temporal correlation of the errors is slightly but noticeable increased. In figure 12 (right) the scatter-plot for the errors of the ensemble forecast for all sites under investigation is represented. For this set the correlation coefficient is increased to about 0.4. It has to be remarked, that due to missing values in the data sets these do not always refer to identical ensembles of sites.

To approve the range of the auto-correlation value for the ensemble the same type of statistical modelling as in equation (2) for the inspection of the standard deviation

12
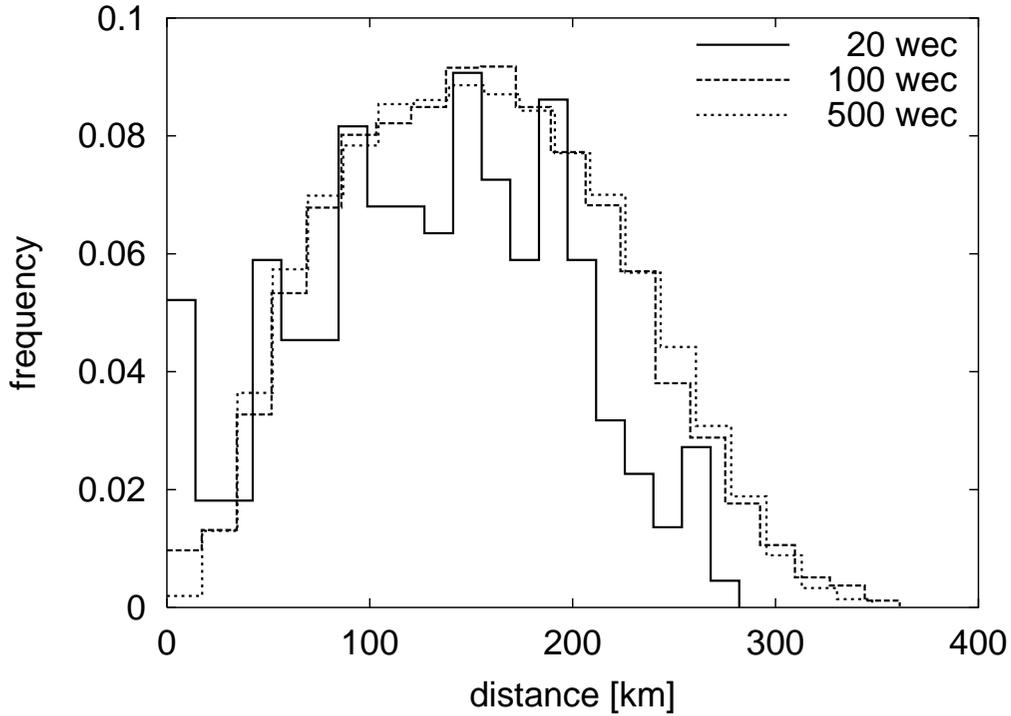
Fig. 10. Frequency distribution of distances in the region with a diameter of 360 km. The curves represent randomly chosen coordinates of 20, 100 and 500 wind farms.

of the forecast errors is used. As key parameters the cross-correlation values of the forecast errors between pairs of sites for a time lag of one day, $r_{xy}(\Delta t)$, has to be taken into account here. The auto-correlation of the forecast error of the ensemble, $r_{\mathrm{ensemble}}(\Delta t)$, is then given by:

$$r_{\mathrm{ensemble}}(\Delta t) = \frac{1}{\sigma^2_{\mathrm{ensemble}}} \frac{1}{N^2} \sum_x \sum_y \sigma_x \sigma_y r_{xy}(\Delta t) \tag{3}$$

For simplicity the temporal average of the forecast errors is assumed to be neglible here. From this information the auto-correlation of the forecast errors with a one day time lag is recalculated for the ensemble of 30 sites. For the 6 h prediction the ensemble error is in good accordance with the respective parameter gained from the scatter plot in figure 12 (right). The dependence of the auto-correlation of the prediction error on the forecast horizon is shown in figure 13 for both the ensemble and a typical single site.

Summing up these findings, it may be stated that, again due to the levelling out of purely stochastic contributions to the single site forecast errors, at least for the 6 h forecasts a somewhat stronger linear link between the ensemble forecast errors of subsequent days exists. However, for the basic ensemble of 30 sites, the respective auto-correlation coefficients are small (below 0.5). As this is still a small value the procedures for an exploitation of this effect have to be refined.

13
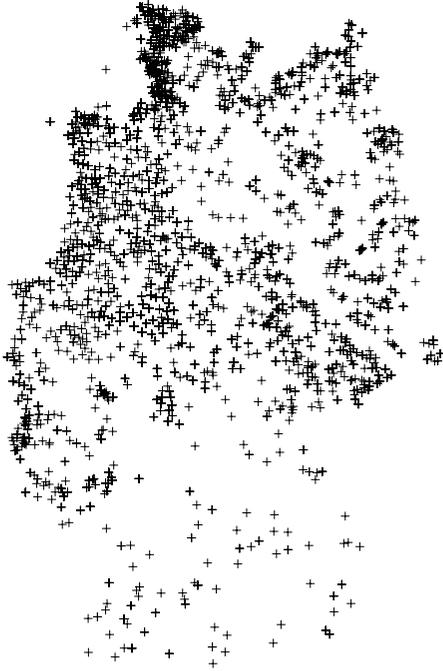
Fig. 11. Distribution of wind turbines in Germany in 1999. The crosses denote ZIP-code areas of the site locations.
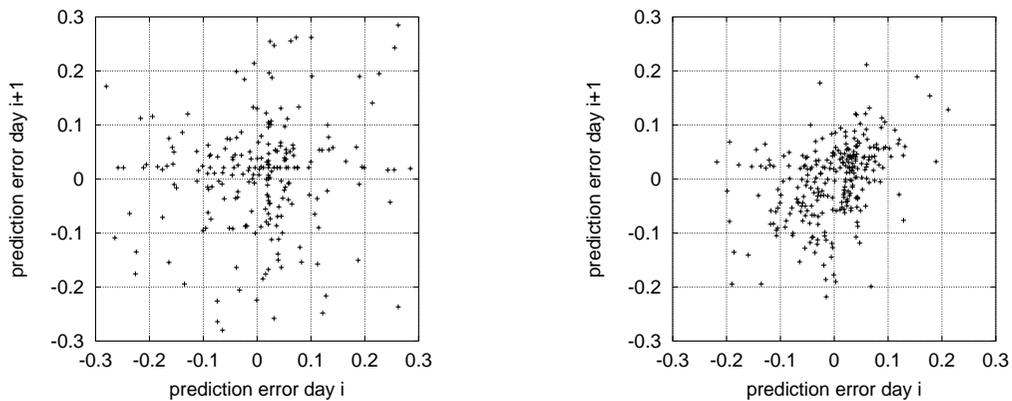
Fig. 12. Left: Example for a scatter diagram of normalized deviations between predicted and measured power for a single site for pairs of 2 consecutive days. Right: Corresponding scatter diagram for normalized ensemble data. Each point refers to the ensemble output of 15-30 sites, depending on data availability.

6 Conclusions

We investigate the statistical smoothing effects of the prediction uncertainty that arise if a wind power prediction is made for a region with spatially distributed wind farms. As expected we find a reduction of the prediction error of the aggregated power prediction compared to a single site. For an ensemble of wind farms where the analysis is based on measured data the improvement of the prediction is no-
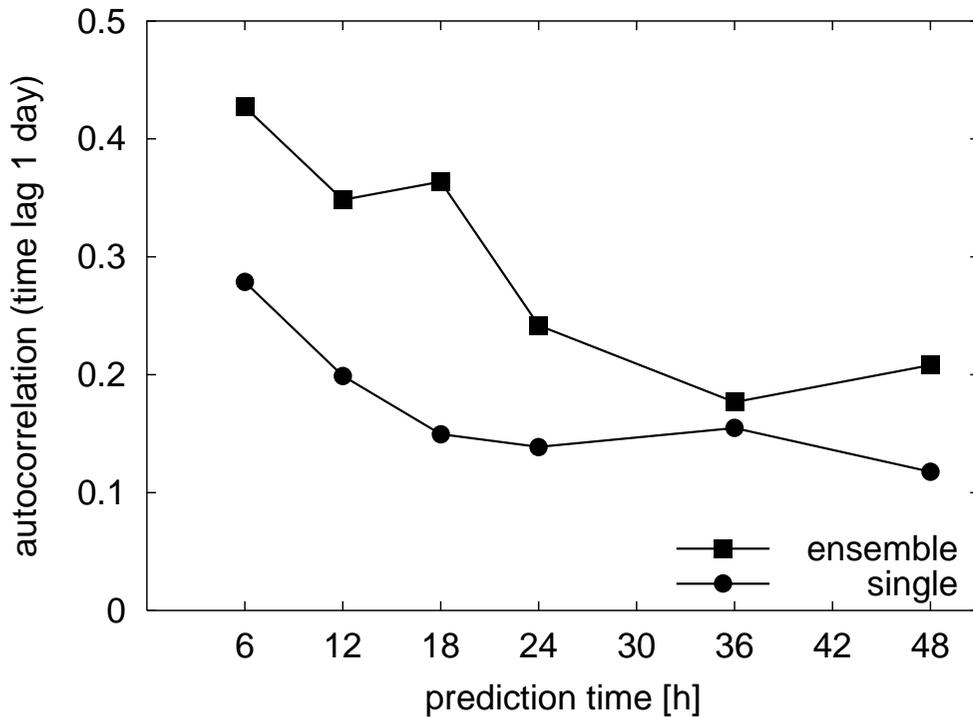
Fig. 13. Auto-correlation coefficients of the prediction error with one day time lag for forecast horizons of 6-48 h. Single site (lower curve) and the ensemble (upper curve). The data for the ensemble are based on pairwise correlation characteristics.

ticeable even for small regions and only few sites. Using model ensembles with randomly chosen locations allows us to generalize the results to identify the impact of the two main parameters, namely the spatial extension of the region and the number of sites it contains together with their distribution. We find that the magnitude of the reduction does strongly depend on the size of region, i.e. the larger the region the larger the reduction. Concerning the number of sites contained in the area we observe a saturation level which is already reached for a comparatively small number of wind farms. This means that only few sites are sufficient to determine the magnitude of the improvement of the power prediction.

The results of our analysis show that for regions with a sufficient number of preferably equally distributed wind farms it is now possible to estimate the regional smoothing effect of the wind power prediction error by just considering the size of the region in question.

As an additional effect of regarding regionally averaged power forecasts a somewhat stronger link between the forecast errors of consecutive days can be identified. This may be beneficially used in the refinement of procedures to correct the actual ensemble power forecast using knowledge on previous errors.

15

Acknowledgements

References

[1]   L. Landberg, Short-term prediction of the power production from wind farms, J. Wind Eng. Ind. Aerodyn. 80 (1999) 207-220.

[2]   H.G. Beyer, D. Heinemann, H. Mellinghoff, K. Mönnich, H.-P. Waldl, Forecast of regional power output of wind turbines, Proceedings of the European Wind Energy Conference EWEC, Nice (1999) 1070-1073.

[3]   K. Mönnich, Vorhersage der Leistungsabgabe netzeinspeisender Windkraftanlagen zur Unterstützung der Kraftwerkseinsatzplanung, Dissertation, Oldenburg (2000).

[4]   T.S. Nielsen, H. Madsen, J. Tofting, Experiences with statistical methods for wind power prediction, Proceedings of the European Wind Energy Conference EWEC, Nice (1999) 1066-1069.

[5]   H. Madsen, T.S. Nielsen, H.A. Nielsen, L. Landberg, G. Giebel, Zephyr and short-term wind power prediction models, Proceedings of the Conference "Wind Energy for the 21st century", Kassel (2000).

[6]   C.S. Moehrlen, E.J. McKeogh, Wind power prediction and power plant scheduling in Ireland, Proceedings of the conference "Wind Energy for the 21st century", Kassel (2000).

[7]   U. Focken, M. Lange, K. Mönnich, H.P. Waldl, Räumliche Ausgleichseffekte bei der Vorhersage der Leistungsabgabe von Windkraftanlagen, Proceedings of the German Wind Energy Conference DEWEK, Wilhelmshaven (2000) 162-165.

[8]   ISET, Wind Energy Report Germany 1999/2000 - Annual Evaluation of WMEP, Kassel (2000).

[9]   H.G. Beyer, J. Luther, R. Steinberger-Willms, Fluctuations in the Combined Power Output from Geographically Distributed Grid Coupled Wind Energy Conversion Systems - An Analysis in the Frequency Domain, Wind Engineering Vol. 14, No. 3 (1990) 179-192

List of Figures

17