On the Benefits of Distributed Generation of Wind Energy in Europe

Vom Fachbereich Physik der
Carl von Ossietzky Universität Oldenburg
zur Erlangung des Grades eines

Doktors der Naturwissenschaften
(Dr. rer. nat.)

angenommene Dissertation

von Gregor Giebel
geb. am 24. 10. 1968 in Wiesbaden

Erstreferent: Prof. Dr. Jürgen Parisi
Koreferent: Prof. Dr. Søren Larsen
Koreferent: Prof. Dr. Jürgen Peinke
Tag der Disputation: 6. 9. 2000

Det er svært at spå, især om fremtiden.
Robert Storm Petersen, gen. Storm P.

Table of Contents

Table of Abbreviations

a.g.l.	above ground level
a.s.l.	above surface level
AT, BE, CH, DE…	Country codes according to ISO 3166 [1]
CC	Capacity Credit
CCGT	Combined Cycle Gas Turbine
CEGB	Central Electricity Generating Board (UK)
CR&W	Combustible Renewables & Waste
DMI	Danish Meteorological Institute
DoE	(US) Department of Energy
DTU	Danish Technical University
DWD	Deutscher Wetterdienst (German meteorological service)
ECMWF	European Centre of Medium-Range Weather Forecasts
EDF	Electricité de France
EdP	Electricidade de Portugal
EIA	Energy Information Authority
EKF	Extended Kalman Filter
ESB	Electricity Supply Board, Ireland
FLH	Full Load Hours
GA	Genetic Algorithms
HIRLAM	HIgh Resolution Limited Area Model
HWP	HIRLAM/WAsP/PARK
IMM	Institute of Mathematical Modelling
IPP	Independent Power Producer
KF	Kalman Filter
KNMI	Koninlijk Nederlands Meteorologisch Institut (Royal Dutch Meteorological Institute)
kWh_{el}	kWh of electrical output
kWh_{th}	kWh of thermal energy
LF	Load Factor
LOLE	Loss Of Load Event
LOLP	Loss Of Load Probability
MAE	Mean Absolute Error
MAPP	Mid American Power Pool
MOS	Model Output Statistics
MSE	Mean Square Error
NCAR	National Centre for Atmospheric Research (US)
NCEP	National Centre for Environmental Protection (US)

NWP	Numerical Weather Prediction
NWS	National Weather Service (US)
OCGT	Open Cycle Gas Turbine
OSTI	Office of Scientific and Technical Information (US DoE)
RAL	Rutherford Appleton Laboratory, UK
RIX	Ruggedness IndeX
RMS	Root Mean Square
RMSE	Root Mean Square Error
SMHI	Sveriges Meteorologiska och Hydrologiska Institut (Swedish Meteorological and Hydrological Institute)
UCPTE	Union pour la Coordination du Transport de l'Electricité
VAR	Variance
WAsP	Wind Atlas Application and Analysis Program
WPPT	Wind Power Prediction Tool

Table of Symbols

$a, b, c, d, \ldots$	Parameters in statistical models
a_k	Parameter in the New Reference Model
A_G, B_G	Stability parameters in the geostrophic drag law
A_W	Weibull magnitude parameter
D	Decay parameter in exponential fit
Eff	Wind farm efficiency
f	Coriolis parameter
$f(u)$	frequency of occurrence of u
$\boldsymbol{F}_i$	State-measurement relationship matrix in the KF
G	Geostrophic wind speed
$\boldsymbol{H}_i$	Autoupdate matrix in the KF
k	Time lag
$\boldsymbol{K_i}$	Kalman gain matrix
k_W	Weibull shape parameter
m_i	Measurement noise in the KF
N	Number of points for analysis
n_i	System noise in the KF
p_k	Measurements at time step k
pen_i	Installed wind capacity in the state model of the capacity credit
q_i	State vector in the KF
s_i	State in the state model of the capacity credit
S	Skill Score
$\boldsymbol{S}_i$	State covariance matrix in the KF
SR1	Fraction of Spinning Reserve depending on actual load
SR2	Fraction of Spinning Reserve depending on wind power
u	Wind speed
u_*	Friction velocity
u_{local}	Predicted wind speed including local effects
$\boldsymbol{V}_i$	Measurement noise covariance matrix in the KF
$\boldsymbol{W}_i$	System noise covariance in the KF
y_i	Observation vector in the KF
z_0	Aerodynamic roughness length
κ	Von Kármán constant
μ	Mean Value
ρ	Multiple correlation coefficient
σ	Standard Deviation

1. Introduction

Wind energy has come of age [2]. As the industry is approaching maturity, the market is shifting from heavily subsidised technology demonstration plant to capital-driven shareholder value, for the industry themselves as well as for the customers. At the same time, the liberalisation of the electricity markets across the OECD has changed the market structure. Instead of the rather clear-cut situation of large vertically integrated utilities with a guaranteed local monopoly, many new players emerge to populate the market.

New players are the independent power producers (IPPs). They also existed in the regulated era, but in those days, IPPs were typically large-scale industry with their own power production facilities, selling surplus electricity with long-term contracts to the regional utility. Now, everyone (in principle) can sell power on the market, to market price conditions. This leads to return-of-investment driven decisions to enter the market as a power producer. To get a reasonable return of investment, knowledge about the market price per kWh attainable is of paramount importance. For IPPs using gas turbines the price during the highest load period is the most important, as this is the main factor entering the economic decision. For investors in a wind farm the average price will be an important factor, or, having multi-year average time series of both price development and wind power production, the price at every hour weighted with the production at this hour. However, the market mechanism is split in two. One is the spot market[a], where the price fixing for the next day happens at noon the day before. The other market is the balance market, where power can be traded on shorter time scales, in case that traders cannot fulfil their obligations. This could be due to a failure of their power plant, or due to an inaccurate prediction of the wind. Typically, the price on the balance market is less than the price on the spot market. Therefore, it is of high importance to predict the wind on this horizon. Wind power prediction for 37 hours ahead is far from trivial, but can be done to a certain extent using numerical weather prediction models. The improvement of this technique by using recursive model output statistics will be the topic of chapter 4.

Since the wind is only partly correlated on large distances, it is desirable to combine the generation of many farms distributed over a large area to decrease the variability of the wind power supply. Note that it is not necessarily the IPP that should be running many wind farms in many countries: the reliability benefit can be reaped even better by another new player in the market, the power broker. Their job is to buy and sell power produced by others on the market. An important subspecies is the green power broker. Trying to match the demand at every hour with the supply is less difficult if the supply is less variable. Also, since the forecast errors are even less correlated than the wind power itself, buying from distributed wind power generators leads to better predictability of the resource on the time scale of the spot market, therefore enabling higher prices. The area and mean spread needed to heighten the firm capacity of wind energy will be assessed in chapter 6.

[a] The terminology and the details are from NordPool. However, other European electricity exchanges are likely to be built similarly.

This market set-up does not lead in itself to a stable supply to the customers. Therefore, for every area there exists a system operator responsible for keeping generation balanced with the demand. Often, this will be the remnant of the old utility, running their existing power plants and doing the scheduling of power plants. It is they, who, at high penetrations of renewable energy, are interested in the output of these plants for the next 8 hours, the time it takes to get a new power plant on-line. As long as penetration and hence, the resulting uncertainty of the renewable supply was in the same ballpark as the uncertainty of the load, conventional scheduling techniques were sufficient. However, since in some areas wind energy delivers 20% or more of the electricity demand (*eg* Navarra in northern Spain, the Jutland part of Denmark or the German land of Schleswig-Holstein)[a], the minimum load can at some times be covered solely from wind energy. Luckily, all Europe is interconnected, so that a higher wind energy production in one region can supplant fossil fuels in another.

The lower variability of wind energy on the European scale has another benefit. Since wind energy is strong especially during winter, in the period with the highest demand, it can replace fossil fuel power plants without affecting the loss-of-load probability. The extent to which this is possible is called the capacity credit of wind energy. Some critics believe that the capacity credit assessment for wind energy has been made obsolete with the liberalisation of the electricity markets [3,4]. However, even when money is the only factor deciding whether to build a new power plant, the investor should be interested to know whether the investment can earn capacity payments for providing firm capacity or not. Equally, since (as we will see later) the capacity credit of wind power decreases and the amount of discarded wind energy increases with installed wind power, an investor would want to know where on the curve he or she is when making the investment in additional wind power. Therefore, the amount of fossil fuel plant replaced by wind energy is assessed on the European scale in chapter 7. Since only one year of wind data was available, the analysis of chapter 7 was set in a long-term context using 34 years of reanalysis data in chapter 8.

[a] The development obviously went faster than the utilities thought. From a German utility assessment on wind energy, 1982: "In the opinion of FRG utilities any discussion on high penetration of wind energy (p>=0.05 of a regional grid) is of no relevance for the foreseeable future; their projection results in an upper bound of a 2.5% contribution of wind energy to the electric energy generation in the year 2000." And: "Utilities find it hard to conceive that sites for *hundreds* of large wind energy conversion systems can be found and that operating licenses can be obtained under the existing legal and regulatory restrictions and against the objections of citizens." (emphasis added) [cited after 111]

2. Models and Methods

2.1 Numerical Weather Prediction

2.1.1 General features of NWP models

Numerical Weather Prediction (NWP) is the simulation of atmospheric processes on a computer with the aim of predicting the future development of the atmosphere based on knowledge of the actual state. This can be done, since most of the physical processes in the atmosphere are known. For a good historical overview, see [5].

The actual state is assessed from measurements from synoptic stations all over the world and from satellites. All this data is used after the data assimilation procedure, which essentially tries to single out wrong data and fill in the gaps between the stations (*eg* over the oceans). Since the quality of the model can only be as good as the description of the initial state, the data assimilation procedure is a rather important step. Unfortunately, there are large areas of the globe where the density of high quality observations is sparse, with no alleviation in sight. To some extent, satellite observations can be used to fill in the blanks, but this is not a universal remedy, since satellites have their own problems. Their advantage is to have a common tool for the observation of weather phenomena over land and sea, a problem is that the observations do include few variables, and some of them only come from over the clouds. However, some additional variables can be estimated from the data available.

From this initial description of the state of the atmosphere, the model calculates the future development. The mathematical formulation leads to a system of non-linear partial differential equations that has no analytical solution. However, a numerical solution can be calculated. For reasons of numerical stability, a solution has to be calculated for every small time step in a grid spanning the model domain with the state variables being calculated at every grid point. Typically, about 20-30 variables are used. The maximum time step possible for a stable solution is dependent on the spatial distance of the grid points and the numerical solver used.

With current computing power, grids of down to a few kilometres spatial distance between the points are possible. Even on this scale some of the atmospheric processes cannot be resolved, like localised thermally induced thunderstorms, draughts through valleys, wind overspeeding over small hills, and all the local influences like buildings, shelter belts, rows of trees and so on. With tricks like the nesting of grids, even better resolutions down to a 100 m are possible, but then only for selected focal points [6]. Some of the atmospheric processes are on length and time scales that will hardly ever be successfully resolved with computational means, like turbulence or the building of precipitation particles. These quantities have to be accounted for by parameterisation at every grid point. Since the parameterised quantities influence the directly calculated quantities, the quality of the forecast is subject to drift with imperfect parameterisations.

Using NWPs the forecast length is still limited nowadays by the calculating power your local meteorological institute has available. In most places 48 hours is currently a typical forecast length for the operational model. However, at the time of our data retrieval it was 36 hours. Many institutes also have coarser models with forecast horizons up to 144 hours.

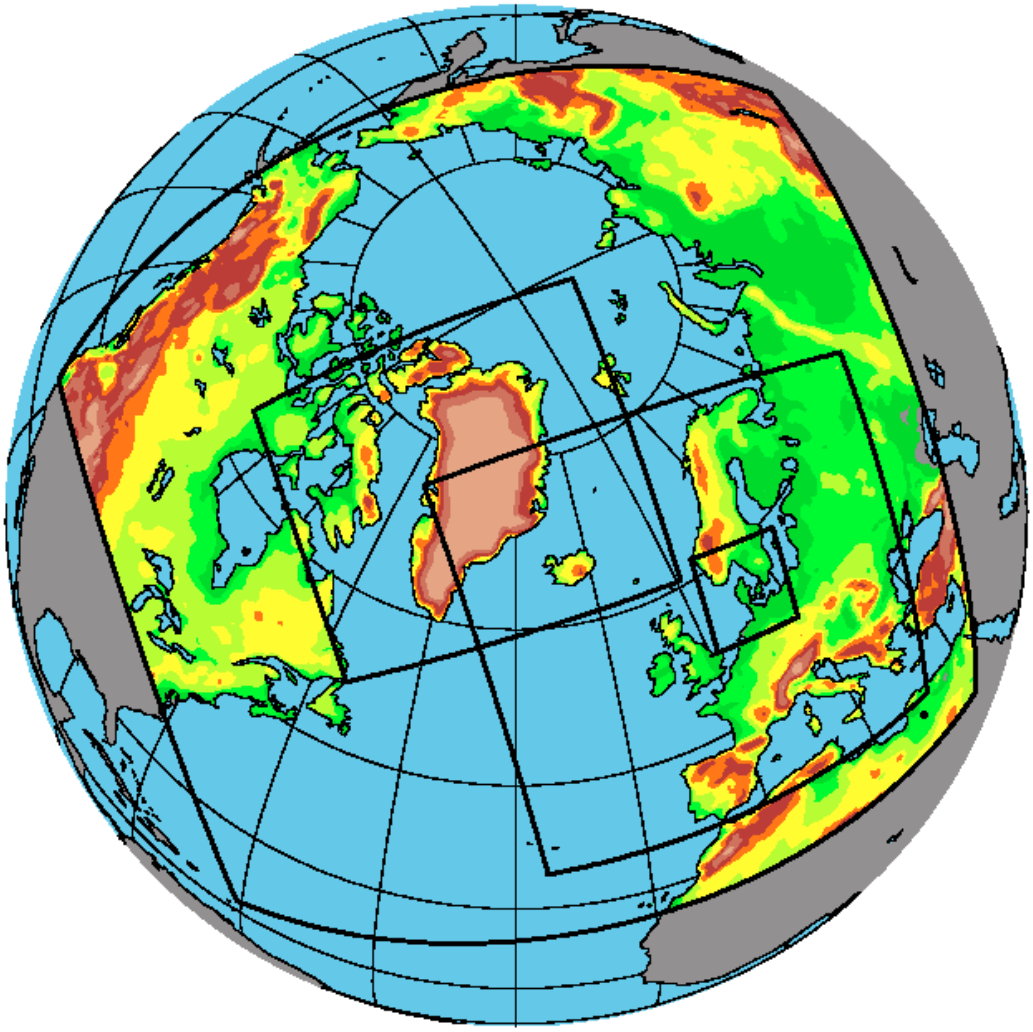

Figure 1: The operational setup of the HIRLAM model.

One theoretical estimate limits the predictability of the weather by NWPs to about 72 hours; this demands, however, that the model domain is global [7]. Lorenz [8] predicts a total divergence of weather patterns from virtually identical starting points after 14-20 days[a], using chaos theory. Using ensemble forecasts, this limit can be extended somewhat, since the ensemble members have some of the possible variation already built in [5].

2.1.2 HIRLAM

The NWP model used mostly in this study is the Danish version of a team effort of the Scandinavian countries, the Netherlands, Ireland and Spain. The HIgh Resolution Limited Area Model (HIRLAM) consists of four submodels using the identical mathematical core, each covering a part of the total domain in various resolutions [9]. The furthest out is HIRLAM-G, which covers an area with cornerpoints in Siberia, California, the Caribbean and Egypt, hence a good share of the Northern Hemisphere. This model is the coarsest, with a horizontal resolution of 48 km and a time step of 240 s.

[a] Nonetheless, the European Centre for Medium Range Weather Forecasts (ECMWF) tries to do seasonal forecasts. This can be done due to the slower time scale of the ocean, which increases predictability to up to 6 months ahead.

The boundary conditions for this model stem from the global ECMWF model [10], which is run twice a day and gives boundary conditions with a 6 hour time step. This model, like the HIRLAM models, has 31 vertical levels. The HIRLAM-G hands over the boundary conditions to two models with a 16km horizontal resolution and 90s time step, one (N) covering Greenland, the other (E) covering Europe. HIRLAM-E is then used to provide the boundary conditions to the model used here, HIRLAM-D, which covers Denmark and parts of northern Germany, western Sweden and southern Norway with a resolution of 5.5 km, with a 30 s time step. The DMI provided us every 12 hours with forecasts in three-hourly time steps for up to 36 hours ahead. These were interpolated from the nearest grid points of the HIRLAM-D to the locations of the farms. The actual set-up can be found in Figure 1.

2.1.3 Nested Grid Model

The National Weather Service (NWS) in the United States provides the Nested Grid Model, which is run twice daily at the National Centre for Environmental Protection (NCEP). The data were provided by the National Centre for Atmospheric Research (NCAR). Among the many models available, the Nested Grid Model was chosen despite its relatively coarse resolution, since it was the only data set complete enough for the purpose of the overall task.

The model uses a sigma vertical co-ordinate system. It has 16 vertical levels up to the 25 hPa level, with resolution comparable to the global spectral model.

The original version of the model included a three-grid configuration. The innermost grid (Grid C) had the highest horizontal resolution (84 km at 45° N). The largest grid (Grid A) covered the entire Northern Hemisphere and had the lowest horizontal resolution (336 km). The original three-nested grid configuration was changed to a two-grid configuration in 1991. Grid C was expanded in all directions to completely encompass the domain of the original Grid B and extend beyond the North Pole. Grid B was expanded to cover the remainder of the hemispheric domain, rendering Grid A unnecessary. The grid spacing is quite large, ranging from 137 km at 20 degrees latitude to 204 km at 90 degrees latitude. At 40 degrees latitude, the distance is 153 km. Forecasts from the 0000 and 1200 UTC cycles are available in gridded form in 6-hour intervals out to 48 hours.

The Nested Grid Model will be dropped from the NCEP production suite when the AVN MOS is developed. The model may be used as a possible component of the short-range ensemble forecasting system.

Nested Grid Model data application of the Risø model to the wind speed and power data bases which were available required access to historical numerical weather prediction data. The data period for the meteorological data from Iowa was June 1994 to April 1997. The data for a 9 x 18 grid encompassing an area from the North Central US to Southern Texas included U- and V-component of the wind at 950, 850, and 700mb, as well as the U- and V-component at 10 meters above ground level. The period of record was January 1, 1994 to March 31, 1996.

2.2 Local scale models

2.2.1 WAsP

The Wind Atlas Application and Analysis Program WAsP [11] has been developed at Risø National Laboratory for the European Wind Atlas [12]. Usually it is used for resource assessment in flat to medium complex terrain, to translate meteorological measurements from one site to another one in the vicinity, taking the local terrain, the local roughness and obstacles around the meteorological site into account. The idea of WAsP is to clean the measurements of local effects, to obtain a wind climate that is representative for the region around the met station, and reintroduce the local effects of the site in question. This is done with three submodels: One describing the effects of hills and other orography features, one describing the effects of different roughnesses on the wind, and a third one describing shelter effects behind obstacles. The resulting wind atlas is described as a set of Weibull A_W and k_W parameters for each sector (see the Appendix for an explanation of the Weibull distribution). Hence, wind speed distributions not fitting Weibull statistics are not well suited for a WAsP analysis. Typically, these occur when a large fraction of the wind comes from driving forces other than the pressure gradient field of the atmosphere, such as thermally induced winds. In addition, the stability treatment of the atmospheric flow is rather generic.

The model for the orography goes back to Jackson and Hunt [13], but has since been updated by several others [14,15,16,17]. Troen and Petersen [18] then moulded it into its current form within their work on the European Wind Atlas. The main feature modelled is the speed-up on top of a hill and the corresponding deceleration in the valley. This special model is based on potential flow, which means that the equations depend on a potential only. This also means that the model does not consider flow separation. Although theoretically well understood [19], the flow separation described in the K-ε theory is numerically more expensive. Therefore, it was chosen to use the simple model. With gentle terrain, flow separation does not occur. The definition of gentle terrain used to describe the operational envelope of WAsP is embodied in the Ruggedness IndeX RIX [20,21]. According to Wood [22], the onset of flow separation is at a slope of 0.3. Hence, the fraction of the surrounding terrain with a slope higher than this critical slope is defined as the RIX. Pairs of stations in areas with similar RIX usually give good resource estimates, since the errors introduced by the orography model are cancelling each other to a certain extent. For stations in areas with different RIX, the error scales linearly with the difference in RIX value.

The roughness model of WAsP is used to generate one effective roughness per sector at a given height, typically the hub height of the turbine. The concept of roughness is linked to the surface stress of the wind, where the wind over a smooth surface is decelerated differently than over a rough surface. At a roughness change, the surface stress changes abruptly, and this change propagates upwardly as the wind moves past the change line. An internal boundary layer develops, where the wind above has not yet seen the effect of the change, while the wind below is already completely in equilibrium. Within the internal boundary layer, a superposition of both effects is taking place. After about 10km, the wind has reached a new equilibrium with the surrounding roughness. For distances from the turbine shorter than this, WAsP takes all roughness changes subsequently into account, according to the theories of Rao *et al* [23] and the measurements of Sempreviva *et al* [24]. The resulting average roughness is strictly valid for the hub

height only, but since the roughness in most landscapes typical for a wind farm does not vary much, the roughness is representative also for other heights. This is not true for sites close to seashore. There, the roughness change is rather pronounced.
The obstacle model of WAsP was not used here, since for any reasonable wind farm there should be no obstacles nearby. Also the meteorological stations used did not have any obstacles nearby.
WAsP compares favourably to similar models, even in difficult terrain [25]. Two recent studies ascertained the accuracy of WAsP in real world situations: Krieg [26] showed that for most wind farms in Sweden the calculated and measured yearly production differed by not more than ±15%. However, it was deemed important to normalise the calculated yearly production to a standard year, since the deviations from the long-term mean wind power output even for 5-yearly means reached over 20%. Mortensen [27] used a derivative of the WAsP methodology for all of Denmark and compared the results to measured production data. More than 80% of all turbines were within 10% of the calculated production.

2.2.2 PARK

The PARK program [28] takes the reduction in wind speed behind the turbines due to wake effects into account. It is therefore used to establish a mean efficiency for any of the turbines in the wind farm. The underlying theory was developed by Jensen [29] and states that the wake spreads linearly behind the turbine. The only parameters going into the model are the initial velocity deficit at the start, and the wake decay constant describing the expansion of the wake. The necessary input for the program is therefore the coordinates of the turbines, the power and thrust curves, the hub height and rotor diameter and meteorological data for the site. The program is limited to wind farms consisting of identical turbines, which in this case was no limitation. The limitation to inter-turbine distances of more than 4 rotor diameters was no problem here, either. The output is one number per sector, giving the efficiency of the wind farm. A comparison of PARK with similar programs has been done by Waldl [30].

2.3 Short Term Forecasting

In the world of short-term forecasting of wind power, three model families can be distinguished according to their input data. One has just the local measurements available, and uses time series analysis techniques with statistical models or neural networks. Another suite of prediction models uses input from NWPs, since for the horizons interesting to utilities persistence and related models do not work well. Finally, the obvious choice for an all-encompassing model is to include both measurements and NWPs. The following section will provide an overview of these approaches.

2.3.1 Persistence and similar models

One of the easiest prediction models (only second to predicting the mean value for all times) is the persistence model, also called the naïve predictor). In this model, the forecast for all times ahead is set to the value now. Hence, by definition the error for zero time steps ahead is zero. For short prediction horizons (e.g., a few minutes or hours), this model is the benchmark all other prediction models have to beat. This is because the time scales in the atmosphere are in the order of days (at least here in Europe). It takes about one to three days for a low-pressure system to cross the continent. Since the pressure

systems are the driving force for the wind, the rest of the atmosphere has time scales of that order. High-pressure systems can be more stationary, but these are typically not associated with high winds, and therefore not so important in this respect. To predict much better than persistence using the same input, that is, online measurements of the predictand, is only possible with some effort.

Bossanyi [31] used a Kalman Filter with the last 6 values as input and got up to 10% improvement in the RMS error over persistence for 1-min averaged data for the prediction of the next time step. This improvement decreased for longer averages, and disappeared completely for 1-hourly averages.

A similar approach is currently used in Wilhelmshaven [32] for the estimation of the wind with the aim of flicker reduction. Vihriälä [33] uses a Kalman filter for the control of a variable speed wind turbine.

Dambrosio and Fortunato [34] used a one-step-ahead adaptive control by means of a recursive least squares algorithm for the electrical part of the turbine. They show a fast and reliable response to a step in the wind.

Nogaret *et al* [35] reported that for the control system of a medium size island system, persistent forecasting is best with an average of the last 2 or 3 values, i.e. 20-30 minutes.

Tantareanu [36] found that ARMA models can perform up to 30% better than persistence for 3-10 steps ahead in 4-sec averages of 2.5Hz-sampled data.

Dutton *et al* [37] used a linear autoregressive model and an adaptive fuzzy logic based model for the cases of Crete and Shetland. They found minor improvements over persistence for a forecasting horizon of 2 hours, but up to 20% in RMS error improvement for 8 hours horizon. However, for longer horizons, the 95% confidence band contained most of the likely wind speed values, and therefore a meteorological-based approach was deemed more promising on this time scale.

In the same team, Kariniotakis *et al* [38, 39] were testing various methods of forecasting for the Greek island of Crete. These included adaptive linear models, adaptive fuzzy logic models and wavelet based models.

2.3.2 Neural networks

Another possibility to use just the input from online measurements is to use artificial neural networks. Most groups in the field have used them, but despite their scientific merits in improvements over plain persistence, they did not catch on. The improvements attainable were usually deemed not enough to warrant the extra effort in training the neural networks. Actually, as we will see in chapter 4.2, most of the improvements found for the persistence-derived methods as well as for the neural networks can be explained fairly easily with the inclusion of a running mean.

Beyer *et al* [40] found improvements in RMS error for next-step forecasting of either 1-min or 10-min averages to be in the range of 10% over persistence. This improvement was achieved with a rather simple topology, while more complex neural network structures did not improve the results further. A limitation was found in extreme events that were not contained in the data set used to train the neural network.

Tande and Landberg [41] examined 10s forecasts for the 1s average output of a wind turbine and found that the neural networks did perform only marginally better than persistence.

Alexiadis *et al* [42] used the differences of wind speeds from their moving averages (differenced pattern method) and found this technique to be superior to the wind speed normally used as input. They

achieved improvements of up to 13% over persistence, while for the same time series the standard neural network approach yielded only 9.5% improvement.
Bechrakis and Sparis [43] used neural networks to utilise information from the upwind direction. Their paper does not give any numbers on the increase over persistence, since their aim is to predict the resource rather than to do short-term prediction.

2.3.3 The Risø Model

Landberg [44] developed a short-term prediction model based on physical reasoning similar to the methodology developed for the European Wind Atlas [12]. The idea is to use the wind speed and direction from a NWP, then transform this wind to the local site, then to use the power curve and finally to modify this with the park efficiency. This general idea is shown in Figure 2. Note that the statistical improvement module MOS can either set in before the transformation to the local wind, or before the transformation to power, or at the end of the model chain trying to change the power. A combination of all these is also possible. Landberg used the Danish or Risø version for all the parts in the model: the HIRLAM model of the DMI as NWP input, the WAsP model from Risø to convert the wind to the local conditions and the Risø PARK model to account for the lower output in a wind park due to wake effects. Two general possibilities for the transformation of the HIRLAM wind to the local conditions exist: the wind could be from one of the higher levels in the atmosphere, and hence be treated as a geostrophic wind, or the wind could be the NWPs offering for the wind in 10m a.g.l. Usually this wind will not be very accurately tailored to the local conditions, but will be a rather general wind over an average roughness representative for the area modelled at the grid point. In the NWP, even orography on a scale smaller than the spatial resolution of the model is frequently parameterised as roughness. If the wind from the upper level is used, the procedure is as follows: from the geostrophic wind and the local roughness, the friction velocity u_* is calculated using the geostrophic drag law (see Appendix). This is then used in the logarithmic height profile, again together with the local roughness. If the wind is already the 10m-wind, then the logarithmic profile can be used directly.
The site assessment regarding roughness is done as input for WAsP. There, either a roughness rose or a roughness map is needed. From this, WAsP determines an average roughness at hub height. This is the roughness used in the geostrophic drag law or the logarithmic profile.[a] In his original work, Landberg [45] determined the ideal HIRLAM level to be modelling level 27, since this gave the best results. However, the DMI changed the operational HIRLAM model in June 1998, and Joensen *et al* [46] found that after the change the 10m-wind was much better than the winds from the higher levels. So in the last iterations of the Risø model, the 10m-wind is used. Since the results here are mostly concerned about the value of MOS to the results, and since the errors associated with these two approaches are qualitatively similar, both were used in the course of this work.

[a] In Oldenburg, the geostrophic profile is used in conjunction with the roughness used by the NWP, not the mesoscale roughness.

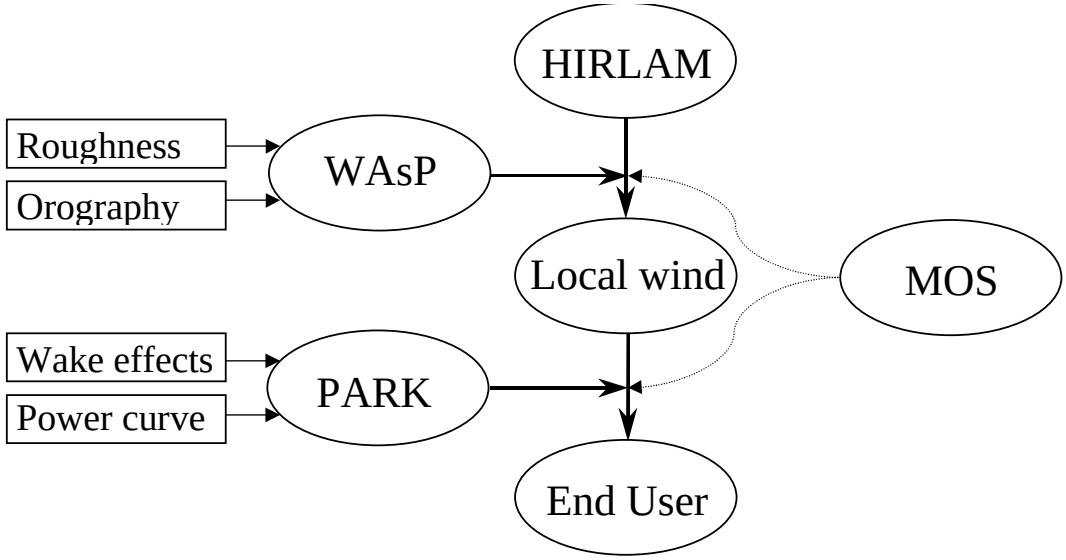

Figure 2: The set-up of the Risø model.

The model runs operatively and is currently used in the dispatch centre of SEAS and Elkraft, the utilities for Sjælland. However, it lacks an upscaling model to relate the wind farm predictions to the output of the whole region.

2.3.4 Other models

The model with the longest operational use with an actual utility is the Wind Power Prediction Tool (WPPT) of the Institute of Mathematical Modelling (IMM) of the Danish Technical University (DTU). The WPPT has been used since July 1994 in the control rooms of Elsam and Eltra, the systems operators in the Jutland/Funen area of Denmark. In the beginning, it only used online measurements of wind speed and total production from seven selected wind farms, which were deemed to be representative of the area they are located in [47]. A data-cleaning module was developed, as was a rudimentary upscaling model. They used adaptive recursive least squares estimation with exponential forgetting in a multi-step set-up to predict from 0.5 up to 36 hours ahead. However, due to the lack of quality in the results for the higher prediction horizons, the forecasts were only used operationally up to 12 hours ahead. In a later version, HIRLAM forecasts were added [48], which allowed the range of useful forecasts to be extended to 39 hours ahead. This version is successfully in operation at Elsam now [49].

A rather similar approach to the Risø model was developed at the University of Oldenburg [50,51]. The main difference here is the use of the Deutschlandmodell of the German Weather Service DWD instead of HIRLAM. A good overview over the parameters and models influencing the result of a meteorological short-term forecasting system has been given by Mönnich [50]. He found that the most important of the various submodels being used is the model for the atmospheric stability. The submodels for orography and roughness were not always able to improve the results. The use of Model Output Statistics was deemed very useful. However, since the NWP model changed frequently, the use of a recursive technique was recommended. A large influence was found regarding the power curve. The theoretical power curve given by the manufacturer and the power curve found from data could be rather different. Actually, even the power curve estimated from data from different years could show strong

differences. The latter might be due to a complete overhaul of the turbine. The largest influence on the error was deemed to come from the NWP model itself. Since the correlation between forecast errors were rather weak with distance, the forecasts for a region were much more accurate than the forecast for single wind farms.

EWind is an US-American model by TrueWind, Inc [52]. Instead of using a once-and-for-all parameterisation for the local effects, like the Risø approach does with WAsP, they run the ForeWind numerical weather model as a meso-scale model using boundary conditions from a regional weather model. This way, more physical processes are captured, and the prediction can be tailored better to the local site. Nonetheless, they use adaptive statistics to iron out the last systematic errors. Their forecast horizon is 48 hours. They just published a 50% improvement in RMSE over persistence in the 12-36 hour range[53].

Vitec AB from Sweden is working on a model based on meteorological forecasts from the Swedish Meteorological and Hydrological Institute SMHI. So far, nothing is published [54].

In a similar setup to the Risø model, just with a different target, Jacobs [55] uses a Kalman Filter to forecast road surface temperatures in the Netherlands based on the 2m temperatures of the HIRLAM model of the Royal Dutch Meteorological Institute (KNMI).

Martin *et al* [56] started to develop a prediction tool for the rather special case of Tarifa/Spain. Due to the unique situation of the wind farms at the Strait of Gibraltar, they could predict the power output from pressure differences between the measurements at Jerez and Malaga airports (west and east of Gibraltar), with the additional use of the Spanish HIRLAM. However, since the utilities felt at that time that 48 hours of forecasts would not be useful enough, the project was stopped half-way through [57].

Papke *et al* [58] used a data assimilation technique together with three models to get a forecast of about 1 hour for the wind fed into the Schleswag grid in the German land of Schleswig-Holstein. These three models were a statistical model, analysing the trend of the last three hours, a translatorical model which moved a measured weather situation over the utility's area, and a meteorological model based on very simple pressure difference calculations. No accuracy was given. The translatorical model developed into the Pelwin system [59]. On a time scale of one hour, the weather fronts coming over the North Sea to Schleswig-Holstein are predicted to predict high negative gradients due to the shutdown of wind turbines.

Another translatorical model was proposed by Alexiadis *et al* [42], which uses a cleaning of local influence much like the methodology used in the European Wind Atlas. The Spatial Correlation Predictor avoids the drawback of the usual constant delay method and shows improvements over the latter of up to 30% and more.

2.4 Model Output Statistics

In any model, especially the mathematically comparatively simple forecasting model[a] of Risø [60,61], systematic errors can occur. This can be due to a wrong roughness assessment, not modelling the effects of atmospheric stability (WAsP uses an average stability, while in reality the actual stability influences

[a] That is, the operational model. The WAsP calculations are more complex, but are done only once.

the scale-up from the 10m-wind to hub height), not modelling the effects of complex terrain, or because the resolution of the NWP calculates the output for grid points in the area of the farm that are qualitatively different. All these effects can lead to more or less complex relations between model output and measurements. In many cases, a simple linear relationship can be established, since the scatter (*ie* the model error) is quite pronounced. Hence, even a simple output filter for the model output can enhance the results significantly. This approach is called Model Output Statistics (MOS) [62].

Generally speaking, there are three things to be considered when using MOS:

- the mathematical relationship between measurements and model output,
- the way in which MOS is implemented, and
- the quality function used.

The relationship between model output and measurements can be very simple or very complex. If there is knowledge available on the mechanisms leading to model errors, this knowledge can be used to parameterise the MOS. However, if this knowledge is not available, a simple linear relationship will in many cases lead to gains in precision. This is also the case if the error mechanisms are too complex and present too many degrees of freedom to be successfully fitted with the data available. Connected to this is the question of whether the MOS should be a part of the modelling process or whether it should be an "end-of-pipe"-technology. In our case it might be better to use MOS on the wind speed (which is an intermediate result) than on the power output.

The way in which MOS is implemented is primarily the decision whether to use static or recursive MOS. With a static implementation, the model parameters are estimated once from older data and then remain fixed. This has the advantage of being very simple and computationally inexpensive. The drawback is that one needs a certain minimum of data to estimate the parameters with reasonable accuracy. Landberg [44] has estimated this to be about 4 months worth of data, which might not be available when setting up the model for a new customer. Another drawback is that varying sets of parameters are not accounted for. There can be many reasons for variations in the parameters, some of which could be parameterised themselves, *eg* the daily or seasonal variation of the wind speed. However, this parameterisation for the seasonal variation ideally needs a long-term data set, spanning at least a few years to be usable. Another change in the parameters can be introduced by changes in the NWP. As has been pointed out in chapter 2.3.3, the HIRLAM model changed in June 1998, rendering all previous MOS parameters useless[a]. To cover for these effects, a recursive model can be used. Recursive means that the parameters are estimated online every time a new model output/measurement pair comes in. Recursive model implementations usually have two steps, one in which the new forecast is produced, and another where the set of parameters is updated. Both steps are independent of each other, since the last estimated set of parameters is always saved. Recursive MOS implementations usually take only a certain number of previous data into account. In order to not run through previous data all over again, the information about past values is stored in a continuously updated covariance (or

[a] A good visualisation of the effect of a change of model came from in climate modelling. There, the average temperature over the ocean jumped quite significantly for no apparent reason. This was later attributed to a model change. This was one of the main reasons to start the reanalysis projects.

similar) matrix. The amount of previous data used is described explicitly in some implementations by a forgetting factor, which is a measure for the effective number of previous data points used for the update. This could also described as the stiffness of the filter - using many previous data points leads to slowly varying parameters, while using only few data points leads to quick variations of the parameters.
The quality function used is of importance for the implementation as well as for the mathematical model. The most common quality function is the Root Mean Square Error (RMSE) (see Appendix for an explanation and formulae of all error functions). Other functions that could be used are the Mean Error (ME), the Mean Absolute Error (MAE), or the skill score. The use of the MAE as a quality function is mainly due to the fact that the overall usability of the forecasts depends more on the general proximity of the forecast to the data than on a few outliers, which are heavily weighted by the RMS. In fact, the monetary penalty for wrong predictions is likely to be directly dependent on the magnitude of the forecast error, since the producer has to make up for the underpredictions by selling the surplus energy rather cheaply on the balance market. This will earn much less than could be achieved on the spot market. In the case of overprediction, the wind farmer has to buy additional kWh. Assuming that the prediction error is not large enough to change the price level on these markets, the financial penalty will only be dependent on the magnitude of the error. That is the reason for the choice of MAE as a quality function. Note that the ME cannot be reasonably used as a quality function since it does not contain information about the overall magnitude of the errors.
These error measures work well when used for the same farm and the same time series. Farms with differently variable time series are not that easy to compare. For this reason the skill score was developed, which takes the different variability of the time series into account. In this way, different results can be compared against each other, without having to worry about the properties of the different time series.
Among the most important forecasts are the forecasts of sudden and pronounced changes, like a storm front passing the utility's area. To develop a measure for the quality of these forecasts is very difficult, however, and the best way to get a feeling for the quality of the forecasts is visual inspection of the data set [*eg* 63]. Other uses of short-term prediction, related to storms, are the possibility of scheduling maintenance after or during a storm, as has happened in Denmark during the hurricane in Dec 1999. The same applies for maintenance on offshore wind farms, where the sea might be too rough to safely access the turbines.

2.5 Kalman Filter / Extended Kalman Filter

The Kalman Filter is an implementation, *ie*, it is the way of applying MOS. The most straightforward way to do this is a least squares fit to your data set, yielding the best possible set of parameters for the model - that is, a set reducing the RMSE to its minimum. Another, more complex, way to find parameters with a merit function different from the RMSE is the use of an optimising method like the ones described in the following section. This has been used extensively in this analysis, due to the possibility of optimising for either RMSE or MAE and the ease of use for optimisation of the Kalman parameters as well.
The Kalman Filter (KF) [64] is a recursive method of implementation. This means that the parameters used in the statistical model do not have to be fixed, but are calculated again newly for each data point

based on the previous data points and the accuracy of the last forecast. Hence, it will adapt the parameters smoothly, should they change over the dataset. The Kalman Filter itself is a linear implementation, which means that only relationships linear in the parameters can be modelled with the standard KF.

In the Kalman Filter, the system is described by the following equations:

The system equation:

$$q_{i+1} = H_i\, q_i + n_i \qquad (1)$$

and the observation equation:

$$y_i = \boldsymbol{F}_i\, q_i + m_i\ . \qquad (2)$$

The symbols refer to the following quantities: the current state is described by the state vector q_i, its covariance matrix $\boldsymbol{S}_i$, the system noise n_i, and the system noise covariance $\boldsymbol{W}_i$. The measurement is described by the observation vector y_i, the measurement noise m_i, and the measurement noise covariance $\boldsymbol{V}_i$. $\boldsymbol{H}_i$ is the matrix relating the state at time step i to the state at time step $i+1$, in the absence of either a driving function of process noise. Since we do not assume any systematic drift of the model parameters q, $\boldsymbol{H}$ is unity in our application. $\boldsymbol{F}_i$ is the matrix relating the state to the measurements. In our case, this is where the HIRLAM predictions come into play.

The prediction algorithm consists of two steps,

the prediction of states:

$$q_{i|i-1} = \boldsymbol{H}_{i-1}\, q_{i-1}\ , \qquad (3)$$

and the prediction of the covariance matrix of states:

$$\boldsymbol{S}_{i|i-1} = \boldsymbol{H}_{i-1}\, \boldsymbol{S}_{i-1}\, \boldsymbol{H}_{i-1}^{T} + \boldsymbol{W}_{i-1}\ . \qquad (4)$$

Independently of that, but usually in the same loop, the internal parameters of the KF are updated:

A new Kalman gain matrix $\boldsymbol{K}_i$ is calculated:

$$\boldsymbol{K}_i = \boldsymbol{S}_{i|i-1}\, \boldsymbol{F}_i^{T} / (\boldsymbol{F}_i\, \boldsymbol{S}_{i|i-1}\, \boldsymbol{F}_i^{T} + \boldsymbol{V}_i) \qquad (5)$$

The state vector estimation is updated:

$$q_i = q_{i|i-1} + \boldsymbol{K}_i (y_i - \boldsymbol{F}_i\, q_{i|i-1}) \qquad (6)$$

The covariance matrix of states is updated:

$$\boldsymbol{S}_i = (\boldsymbol{1} - \boldsymbol{K}_i\, \boldsymbol{F}_i)\, \boldsymbol{S}_{i|i-1} \qquad (7)$$

A central point is the Kalman gain matrix, which can be written as:

$$\boldsymbol{K}_i = \boldsymbol{S}_i\, \boldsymbol{F}_i^{T} / \boldsymbol{V}_i. \qquad (8)$$

Thus, the gain matrix is proportional to the uncertainty in the estimate and inversely proportional to that in the measurement. The Kalman gain matrix is a weighting factor to bring the KF state estimate back on track when errors occur. If the measurement is uncertain while the state estimate is relatively precise, then the residual stems mainly from noise, and the correction of the state estimate should be small. If on the other hand the uncertainty in the measurement is small and that in the state estimate is large, then the

residual is based on real deviations of the state estimate from the measurements and the state estimate should be corrected strongly accordingly. The Kalman gain matrix, and hence the relative uncertainties in the system noise and the measurement noise, determines the "stiffness" of the filter. Since the system and measurement noise in my study are seen as parameters, the stiffness is determined externally. The noise terms and hence, the stiffness are open to optimisation. These terms also define a maximum step width (in conjunction with the given scatter of the time series to be analysed).
The initialisation of the KF is rather simple, since the initial parameters should not have a great influence on the final estimate, once the KF is in equilibrium.

$$\boldsymbol{S}_0 = \boldsymbol{V}_0 \text{ ; and } q_0 = E[q_0] \,. \qquad (9)$$

The standard KF can only cope with dependencies that are linear in the parameters. To expand the operational envelope of the KF, the Extended Kalman Filter (EKF) can be used for models that are non-linear in the parameters. This filter uses for the internal adjustment a linearisation around the current point in the parameter space. Therefore, it also can be used to adjust the wind before applying the (non-linear) power curve to it. The linearisation however requires a fairly smooth function to work with.[a]
EKF System equation:

$$q_{i+1} = H(q_i) + n_i \,. \qquad (10)$$

EKF Observation equation:

$$y_i = F(q_i) + m_i \,. \qquad (11)$$

This equation will be filled later with the models 1-5 described in section 4.3.
Since the prediction and update steps are very much along the lines of the standard KF, except for the inclusion of a linearisation step, the equations will not be shown here.

2.6 Optimisation: Downhill Simplex, Genetic Algorithms

The aim of an optimisation routine is to find the set of parameters for which the quality function reaches the minimum. This function to be analysed can be rather complex, but as long as it depends linearly on the parameters, a Least Squares Fit of the data will yield the set with the minimum RMSE. However, for the other quality functions and other algorithms a proper optimisation must be used.
Two rather different methods have been chosen here for the optimisation: the Downhill Simplex algorithm by Nelder and Mead [65, very well explained in 66], and a genetic algorithm [67].
The downhill simplex is also called Amoeba [68], due to its creature-like behaviour. It is to be thought of as the (N+1) points in the N dimensions of the problem that enclose an N-dimensional volume. In 2 dimensions, this would be a triangle, in three dimensions a tetrahedron etc. The simplex can contract or expand, and can move by mirroring points on its own surfaces. This mirroring is done for the point with the least merit, therefore the simplex is moving towards a minimum. The starting simplex has to be set

[a] Since the power curve is smooth enough to satisfy the requirements, the linearisation was not a problem. This drawback can be overcome by the use of a more recent technique, the Unscented Filter: Julier, S.J., and J.K. Uhlmann, *A General Method for Approximating Nonlinear Transformations of Probability Distributions*, Internet Publication: *http://www.robots.ox.ac.uk/~siju/work/publications/ A4_size/Unscented.zip*

according to the expectations of where a minimum should occur, but on reasonably smooth surfaces it should not matter from which point the algorithm is to start. The big advantage of this algorithm over other multi-dimensional optimisation algorithms is that there is no need for differentiation, and the algorithm is so general that setting it up to do some work is relatively straightforward. A disadvantage is that only part of the parameter space may be analysed, and the simplex can be fooled by a local minimum. While this holds true for any multi-dimensional minimisation routine, the genetic algorithm described in the next paragraph does a broader search, albeit with the cost of much more computation.

Genetic Algorithms (GA) make use of the evolutionary theory for finding the optimal or sub-optimal solution. Just like the downhill simplex they operate directly on the function, without recourse to its derivatives, they search from the population, but not from a single point, and the fitness function can be any function of the chromosomes. A chromosome is the coding of a solution to be used together with the fitness function. This can *eg* be a parameter vector using a least squares assessment as a fitness function. Initially, the GA will randomly generate a set of chromosomes and form a population. This population will undergo the 3 standard genetic operators: Crossover, Mutation and Selection. The chromosomes inside the population will become better and better. Finally, the algorithm will terminate according to the stopping criteria.

2.7 National Grid Model

The assessment of the economic value of wind energy and wind power forecasting is routinely done at the Rutherford Appleton Laboratory using the National Grid Model (NGM) [69,70]. Detailed information is needed on the power plants available for dispatch, and the prices for fossil fuel. Additionally, three time series are needed in the resolution of the model time step: electricity demand on the whole grid, wind power measurements and possibly wind power forecasts.[a] This tool has been used and improved continually over more than ten years [*eg* 71].

It works by modelling the scheduling and dispatch of power plant to meet the demand on a large-scale electricity grid. As such, it is a one-node model, where all generation is done at the dispatch centre, and all the load is served in one big block from there. This means that transmission losses or bottlenecks are not an issue for the NGM. However, for most of Europe, where nearly all countries are part of UCPTE, this is not far from the truth.

In this set-up, the model runs in hourly time steps. At every step, the number of plants needed in the next hours to cover the predicted demand is scheduled ahead. All power plants are listed in a merit order, which means that the operators choose the plant with the lowest production cost per kWh first, and then work their way through the list up to the number of plant which are able to cover the predicted demand. At this stage, the predicted wind power is treated as negative load. Due to this merit order approach, wind power replaces the least efficient and probably also most polluting power plants first. To account for the uncertainty of the demand, the predicted demand is calculated by multiplying the actual demand with a Gaussian distributed random number with a distribution mean of 1 and a standard

[a] Persistence and perfect forecasting are calculated within the program itself.

deviation of 0.015. This uncertainty is consistent with published deviations for load prediction algorithms [*eg* 72,73,74]. The eight plant types treated in this model can be found in Table 1.

Table 1: Details of the power stations used in the NGM. OCGT=Open Cycle Gas Turbine, CCGT=Combined Cycle Gas Turbine, CR&W=Combustible Renewables & Waste.

*: The cost for OCGT is €c/kWh$_{el}$, all other costs are fuel prices in €c/kWh$_{th}$.

**: Average lifetime load factor

Type	Startup time [h]	Cost [€c/kWh]	Mean efficiency
Nuclear	n.a.	n.a.	0.70**
Coal	8	0.7294	0.35814
Oil	8	0.9489	0.35755
OCGT	0	8.6245*	0.1987
CCGT	1	1.0204	0.54685
CR&W	8	.8793	0.305
Hydro	n.a.	n.a.	n.a.
Pump storage	n.a.	n.a.	0.88475

An assumption is made for each type of plant regarding its start-up time: a maximum of eight hours is assigned to coal- and oil-fired plant, while OCGTs are considered to start up immediately within the time frame of the model. Obviously, the fraction of OCGTs in the modelled grid strongly influences the response time of the grid. The eight-hour maximum also limits the time frame for looking ahead - there is no need to look beyond the maximum start-up time. Nuclear plant is treated as always running with an output determined by the rated capacity times the lifetime load factor. Some additional details will be discussed in chapter 7.5. Hydro plant is predominantly (90%) used as base plant with a generation according to season: in late winter the generation is higher than in late summer. Any shortfalls in load not covered by the scheduled power plant are met by either fast response plant (pumped hydro or gas turbines) or through the spinning reserve.

The spinning reserve is thermal plant, which is not being run at full output, but at, say, 95%. The remaining 5% can be activated very fast if need be. Thermal power plants cannot be operated at less than 50% load factor. Therefore, this is set as the minimum load factor. The spinning reserve is planned as a fraction of the predicted load (SR1) as well as a fraction of currently available wind power (SR2). Both these fractions remain fixed for a model run (typically one year), but are optimised to yield minimum fossil fuel cost under the condition that no loss-of-load-events (LOLE) occur. The condition that no LOLE may occur can lead to a rather high SR2 and hence a high overall spinning reserve requirement. Since power plants can only be dropped from service from one time step to the next in the model, not all of the wind power production can be accepted into the grid when all running steam plants are already at the minimum load or all steam plant are shut off and the demand is already covered by nuclear, hydropower and a fraction of the available wind energy. This means that high values of SR2 at high penetrations of wind energy can also lead to significant wind power production being discarded.

This model set-up is then run for added capacities of 0, 1, 2, 3, 5, 10, 20, 30,…100% of wind energy, as a percentage of the existing plant mix.[a] This percentage refers to added nameplate wind generating capacity in relation to the existing conventional power plant capacity. This is slightly different from penetration, which is the percentage of wind energy in the then bigger grid. Note that the plant mix is treated as steady state, *ie* no plant is built or decommissioned during this year.

The optimising can occur in one of two ways: the more often used (since it is numerically more benign) is to optimise SR1 at 0 additionally built wind energy, and then go through steps with additional wind energy, keeping SR1 at the found value and optimising for SR2 alone. In this thesis, all graphs refer to this set-up. The other option is to optimise for SR1, SR2 and one parameter detailing the strategy for the pump storage plants all at once for all penetrations. However, since the results are very similar to the ones shown here, and since the optimisation occasionally did not converge properly, these graphs were omitted.

To assess the influence of wind power forecasting, two different forecasting methods are used: one is perfect forecasting, and uses the time series as a forecast. The inclusion of this model yields an upper boundary for the usefulness of wind energy forecasting, as well as for wind power in general. The other model is persistence, which is quite reasonable, since the maximum forecast horizon needed is 8 hours. This model sets the lower bound of what is possible with short-term forecasting. Actually, forecasts on the basis of NWPs alone might fare even worse, since these have significant error for the zero hour forecast (see *eg* Figure 13). Nevertheless, in a real world application it would be very unwise not to use online measurement data, and hence to get the zero hour prediction error down.

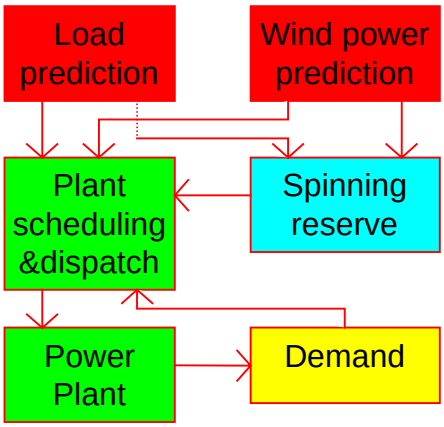

Figure 3: The workings of the NGM.

[a] Actually, the European plant mix was redefined from its first assumptions in a later stage in the project, but the newly installed capacities were left untouched. Therefore, the added capacities are multiples of 0.79%, rather than 1%.

3. Data

3.1 Danish Wind Data

The data used for the analysis of MOS in chapter 4 was hourly production data from 12 wind farms in eastern Denmark and the associated three-hourly HIRLAM forecasts. For a detailed analysis, the Nøjsomhedsodde wind farm on Sjælland was picked. This farm contains 23 Vestas 225-kW-turbines. The data were recorded from February 1995 to August 1998. Since not all turbines were online at all times, the recorded data were divided by the number of running turbines and then scaled back to the number of turbines in the farm, thus giving the behaviour of an "average" turbine representative of the wind farm.

3.2 England and Wales Wind Data

The wind speed data from England and Wales was used for the Kalman Filter results in chapter 4 and for the sensitivity analysis of the NGM in chapter 5. Eleven sites monitored as potential wind farm sites were used as supplied by the UK wind farm developer Renewable Energy Systems Ltd. The data from these sites, which were used for this study, covered the calendar year 1994.

3.3 CEGB England and Wales Grid Data

The database of the England and Wales grid was used for the NGM sensitivity analysis in chapter 5 and for the European grid described in section 3.7. It covered 109 units with a combined peak capacity of 58593 MW. This includes the connectors to France and Scotland, which were modelled as coal fired plant. Hourly load profiles were available from 1971 to 1996. The peak load in 1994 was 47266 MW; the cumulative load amounted to 272.05 TWh.

3.4 Iowa Wind Data

The NGM sensitivity analysis performed in chapter 5 was mainly performed on data from Iowa. In an aggressive program to map out the wind resource in the state, the State of Iowa set up a large monitoring program comprising, amongst other things, of a dozen 50 meter wind monitoring masts throughout the state. The data of two of these, from Alta and Sibley in the north west corner of the state, not far from the better known Buffalo Ridge site, was used in this comparison. The data extended from January, 1994 to March, 1996.

3.5 Iowa Grid Data

The load curve we had for disposal was from a smaller utility within Iowa. An analysis of data available on the webserver of the MAPP [75] (Mid American Power Pool) made it possible to scale the load curve up to yield a load curve representative for all Iowa.

The Iowa grid contained 42 power plants with a combined rating of 7351 MW. The minimum/maximum load in 1995 was 1634/6969 MW, compared to 1564/6172 MW in 1996. The cumulative demand for both years was 347.48 and 347.35 GWh, respectively. Please note that the load data sets in the Iowa case refer to calendar years, while for the England and Wales grid they refer to a

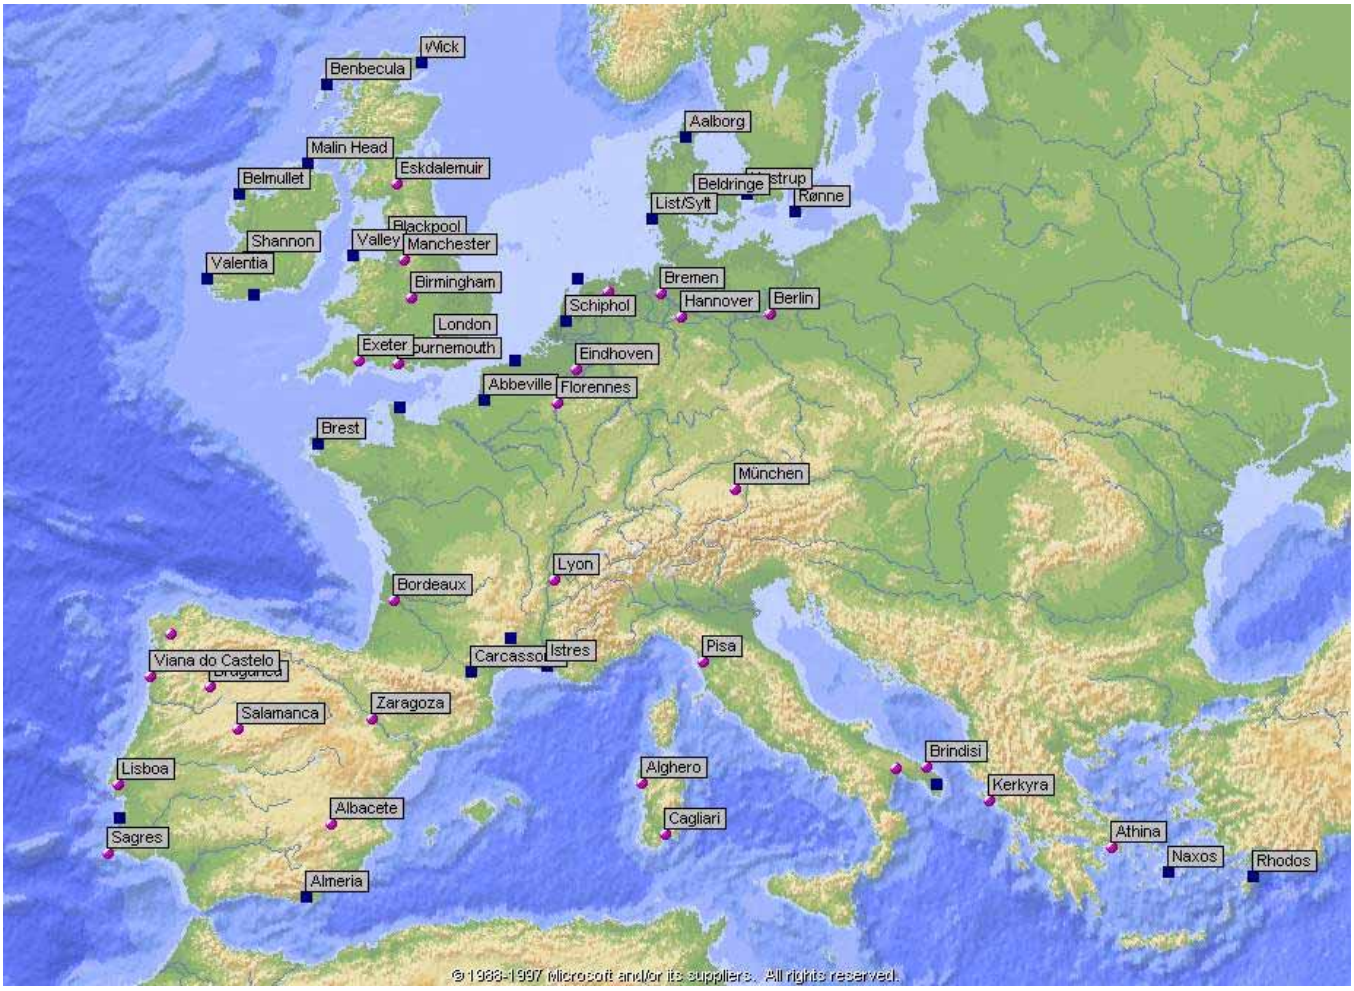

Figure 4: The meteorological sites used. The blue squares are stations used for the selection series. If the name is missing, it means that too much data was missing to use the station on its own.

financial year, which runs from April, 1st, to the end of March. Hence, results denoted 88 in this case are valid for the financial year 1988/89.

3.6 European Wind Data

The European wind data for the analysis in chapter 6 came from 60 meteorological stations in the selected countries and is detailed elsewhere [76,77]. The geographical distribution can be found in Figure 4. The data extended from December 1990 to December 1991. The data quality can be judged from Figure 5. Most data sets were usable, only 8 had nearly all data points missing. These stations were also used in the averaging, but not in analyses carried out for single sites (like the crosscorrelation analysis in chapter 6.1).

Since the time series is only three-hourly, while the NGM needs an hourly time series, the wind speed was linearly interpolated at every station before applying the power curve. The wind was scaled to a height of 50m above ground level, using the sector dependent roughnesses from the WAsP analyses published in [44] in the logarithmic wind profile. In order to calculate the total European wind power generation from these sites, a European average wind turbine distribution was used. The distribution can be found in Table 2. The total power curve is shown in Figure 6 and is a superposition of the power curves of:

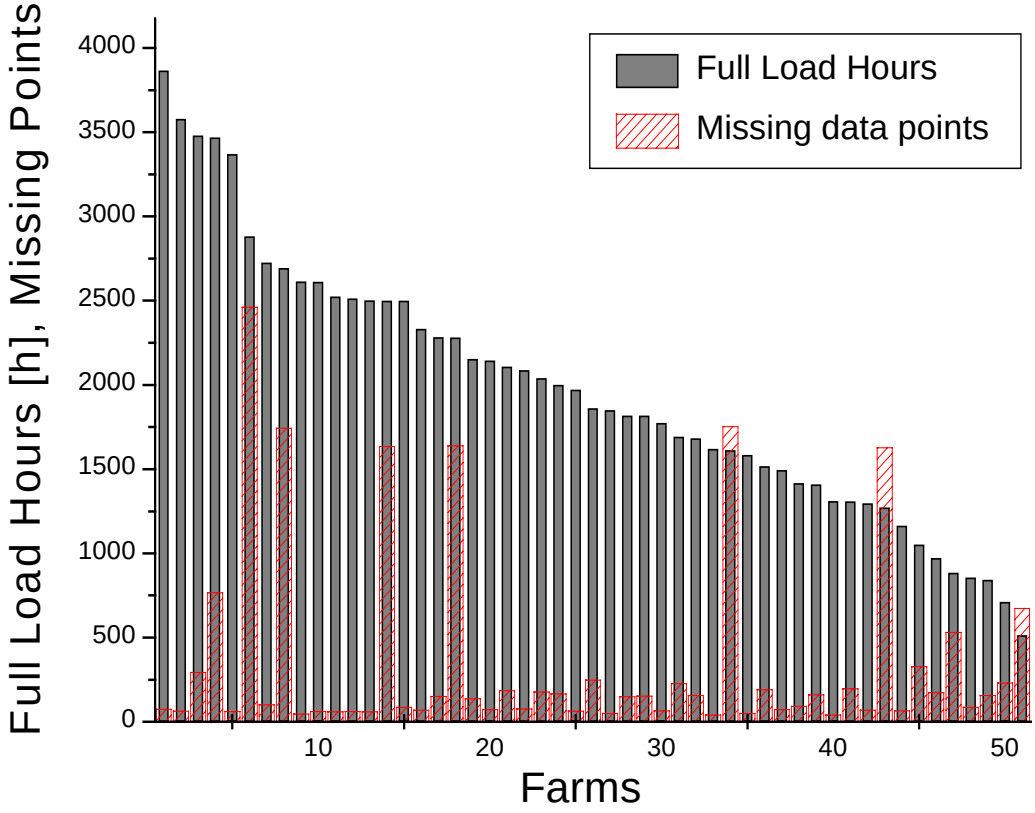

Figure 5: Number of full load hours and number of missing data points in the single time series. Sites missing more than 2500 points were left out in this graph. Here many stations are inland and in practice would only see development where local topographical effects enhance the resource.[a]

Table 2: Overview of wind turbines used to model the European wind turbine distribution.

1	Vestas V66	1650 kW
1	Avedøre test turbine	1000 kW
2	Micon	750 kW
1	Wind World W-3700	500 kW
1	Windane 34	400 kW
2	Vestas V27	225 kW
2	Danwin 27	225 kW
1	Nordtank	150 kW

The sum is 6100 kW; the average is 554.5 kW. This is adequate since the average among newly installed turbines in Germany up until October 1998 was 764 kW, while the installed base rated capacity

[a] An example can be seen when entering Munich from the north. Here, despite the usually rather poor resource, an Enercon 1.5MW machine was erected on a garbage mound, using a 100m tower. Thereby, the local resource was just about good enough to warrant an investment.

was 444 kW/unit [78]. Extrapolating these trends, this turbine distribution should be representative for early 2000.[a]

Using the superposed power curve for each site, the power output time series was aggregated over Europe for each hour. A data point was only used if at least 25 sites had a non-missing wind speed value - otherwise, linear interpolation of the resulting time series was used. This was necessary in 76 cases. This time series is referred to as 'Average'. In order to measure the effects of time series with higher load factors, but also higher variability, two additional series were created: the one called 'Selection' is averaged over the 25 farms with more than 2000 Full Load Hours (FLH), while the series called 'Malin Head' is the single site time series with the most FLH, which came from Malin Head in the Republic of Ireland, with 3865 FLH.

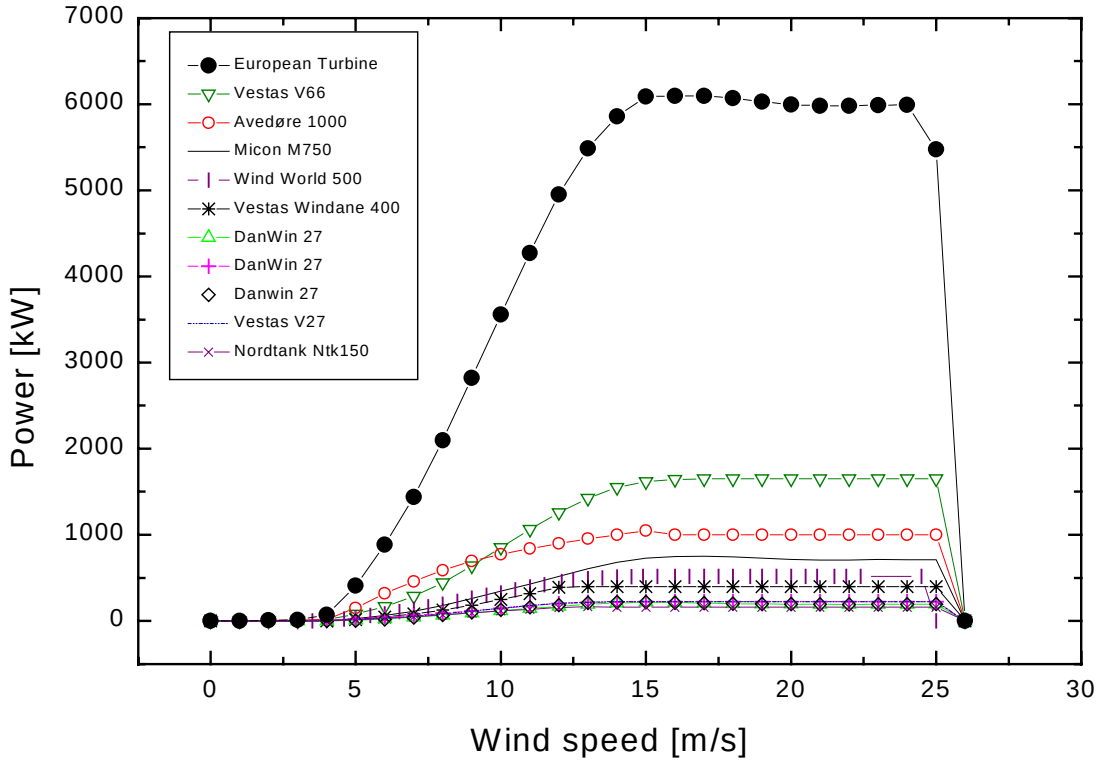

Figure 6: The power curve used to model the European power curve distribution.

3.7 European Grid Data

In order to simulate the European grid (see chapter 7), the details of every power station in Europe, the fuel prices, a full demand time series of the selected countries and the corresponding wind power time series would be needed. Unfortunately, not all of this was available.

The installed capacity in the selected countries (Austria, Belgium, Denmark, France, Germany, Greece, Ireland, Italy, The Netherlands, Portugal, Spain, Switzerland, and the United Kingdom) was available [79], broken down by plant type. Additionally, the full individual power unit details for England/Wales, Ireland and Portugal were known, as were the details of all European nuclear power stations [80]. In order to estimate the distribution of the individual power units for the remaining countries, the known power units were divided into the 8 categories used by the NGM. For each category, the number of units

[a] No newer data was available at REISI at the time of writing.

was scaled up to the appropriate total capacity for the European countries selected. The overall capacity for all categories is 461.42 GW [79]. However, the available capacity at the day of the highest demand was only 321.13 GW. Since the National Grid Model does not take scheduled or forced outages into account, this installed capacity was chosen for the modelling of the whole power plant mix. With the hydro generation assumptions and the nuclear plant-scheduling algorithm employed by the NGM, the installed capacity rose therefore to 347.375 GW.

The demand time series were available from France (EDF), England/Wales (CEGB) and Portugal (EdP). These were scaled in order to fit the overall European load, which was 1603 TWh. Every time series had a weight of 1/3, as determined by the cumulative load in that period. The Portuguese time series was actually a time series for 1992. To get it corrected in respect to the distribution of weekdays and holidays in 1991, whenever possible the more fitting day of the next or previous week was taken. At some days it was necessary to use the original holiday period (especially the Easter weekend) and do a slight linear interpolation of the data to make the ends meet with the previous time series.

3.8 European Wind Data from Reanalysis

The reanalysis project is an effort to reanalyse historical meteorological data with state-of-the-art weather models. The aim here is to create one consistent data set without artificial trends that were introduced at a model change. The meteorological model used is along the lines of the operational model described in chapter 2.1.

In the US effort [81], most meteorological observations from 1948 to present day were used to create one consistent model run, using the same data assimilation procedure and the same meteorological model throughout. Thereby, one consistent time series spanning the whole globe for five decades was created. The model used a T62 spectral triangular Gaussian grid with a resolution of about 2°x2° horizontally and 28 levels vertically.

The data used in chapter 8 was from 1965 to 1998, 12-hourly 10-minute averages of wind speed at 10 m height a.g.l. and various model levels. Unfortunately, the 12-hourly resolution is not good enough to run the data with the National Grid Model. However, for getting an idea of the distribution of winds in some special period over many years, this data is very useful.

The time series were from the grid points nearest to the meteorological stations described in chapter 3.6. Since the mean of the local time series were partly different to the mean of the meteorological data series, the reanalysis data wind speed was renormalised to yield the same mean wind speed as the meteorological data for the period Dec 90 to Dec 91. From there, the conversion to wind power was done using the same methods as the short-term forecasting in chapter 2.3.3. It was tried to use the wind speeds from the 850-hPa-level of the reanalysis data. Since this performed much worse in comparison the 10-m wind speed, the latter is used throughout.

A comparison of the distributions of the measured data and reanalysis derived data (including local effects analog to Risøs forecasting model, chapter 2.3.3) is found in Figure 7. While in general the distributions are fairly similar, in cases with strong local effects the Weibull shape (see Appendix) for the measured wind speeds disappears. Roches Point/IE is such a place, with water on one side and relatively high roughness and hills on the other. Therefore, the Weibull distributions for the measurements are still valid for wind speeds from a certain direction, but when plotting the frequency

distribution of the whole time series, the different Weibull parameters do not add up to a Weibull shape for the whole. Another effect smoothing the statistics in the reanalysis case is that the time series was only 13 months in the case of the meteorological measurements, while the reanalysis series is an average over 34 years.

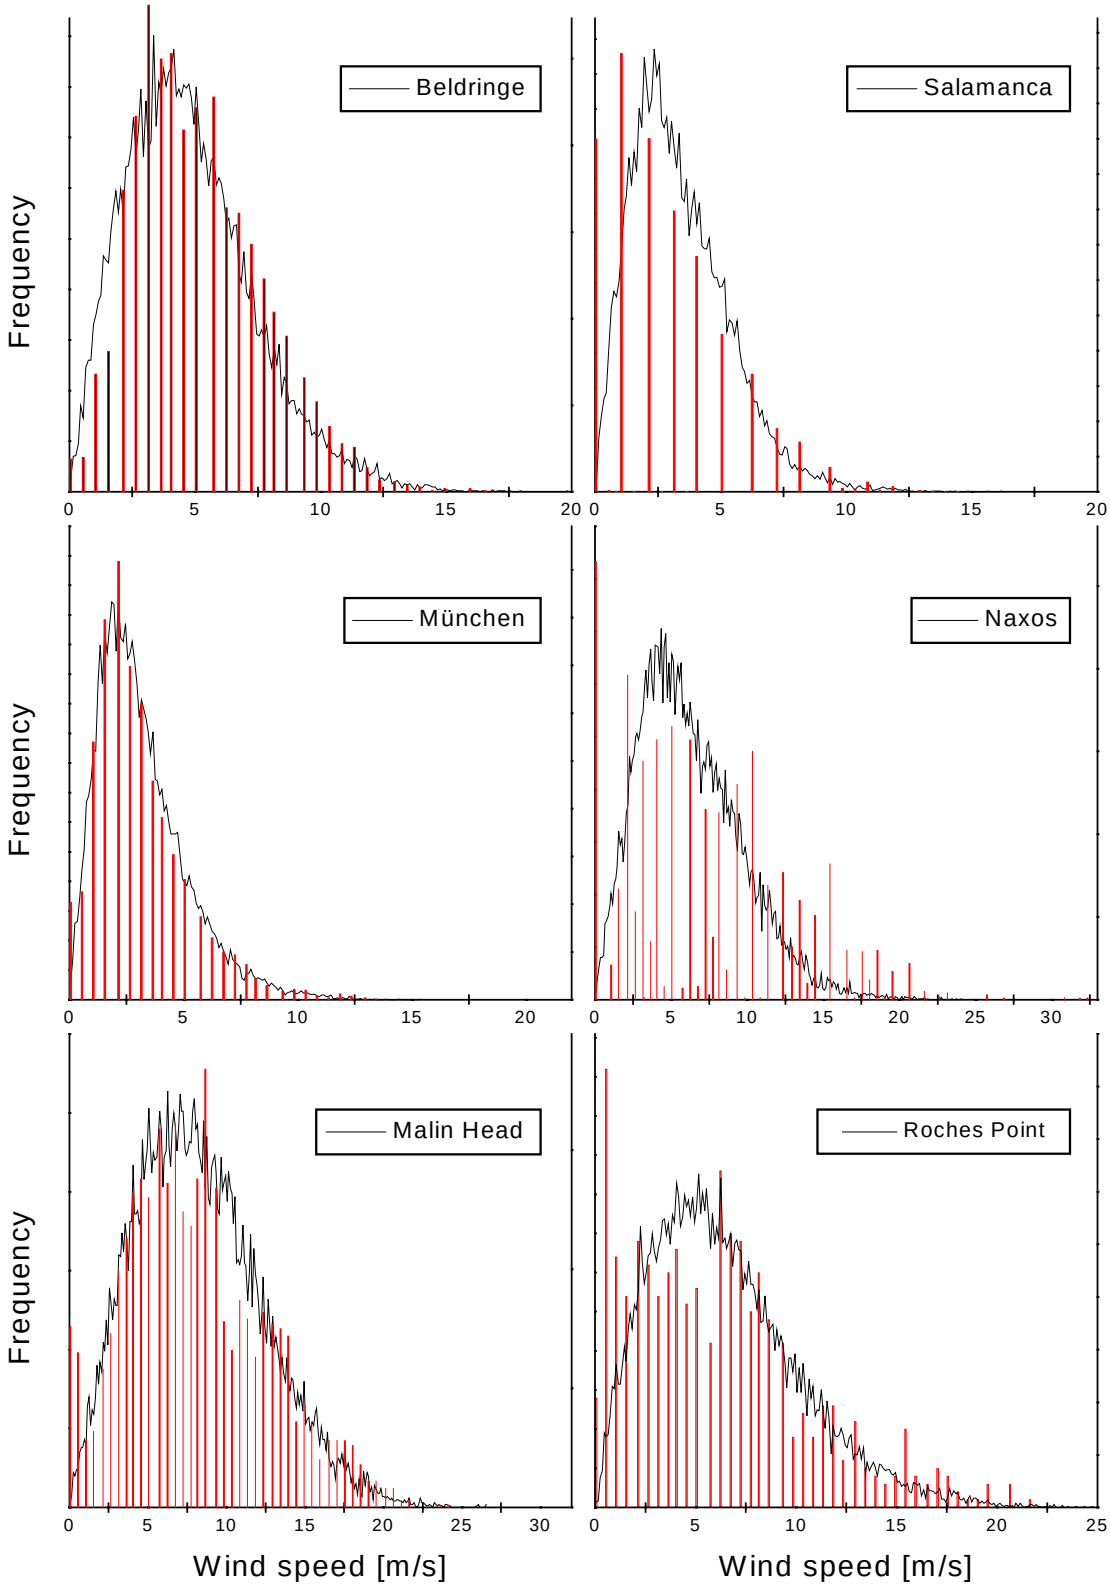

Figure 7: Wind speed distributions for 6 selected sites. Columns are meteorological measurements; full lines come from reanalysis. The period is in both cases the whole data set. Please note that the x-axis is not the same for all plots.

4. Model output statistics

4.1 Introduction

The chain of models introduced in chapter 2.3.3 (the Risø model, using HIRLAM/WAsP/PARK (HWP)) uses physical reasoning, but leaves many effects unmodelled. Among these are the seasonal variation, the daily variation, variations introduced by stability effects, effects of orography or roughness in large wind farms (since the calculation is for one spot only per farm), the uncertainty in the power curve (which can easily be more than 10%), and non-attached flow in complex terrain. Some of these processes, especially the seasonal and daily variation, as well as parts of the stability, should be handled by the NWP. However, since the NWP has limited resolution and calculates with an average roughness for every grid point, many local effects cannot be resolved by the weather model. In these cases, an additional correction term based on empirical parameters can be used. This correction is Model Output Statistics (MOS) [62].

In our case, where some of the errors are varying over the year, while others are not changing, an adaptive solution of error removal would be best. On the other hand, a simple least squares fit to older data already might lead to quite reasonable results. In Figure 8 the error in power and in wind speed is plotted with its frequency. For the wind, a pure Gaussian with a σ of (in this case) 1.92 m/s can be fitted quite satisfactorily, while for the power the error is a superposition of two Gaussians with σ of 19 and 602 kW, respectively. The high frequency of near zero error comes from two parts of the power curve: the flat part before the cut-in speed, and the flat part at full rated power. There, the wind speed

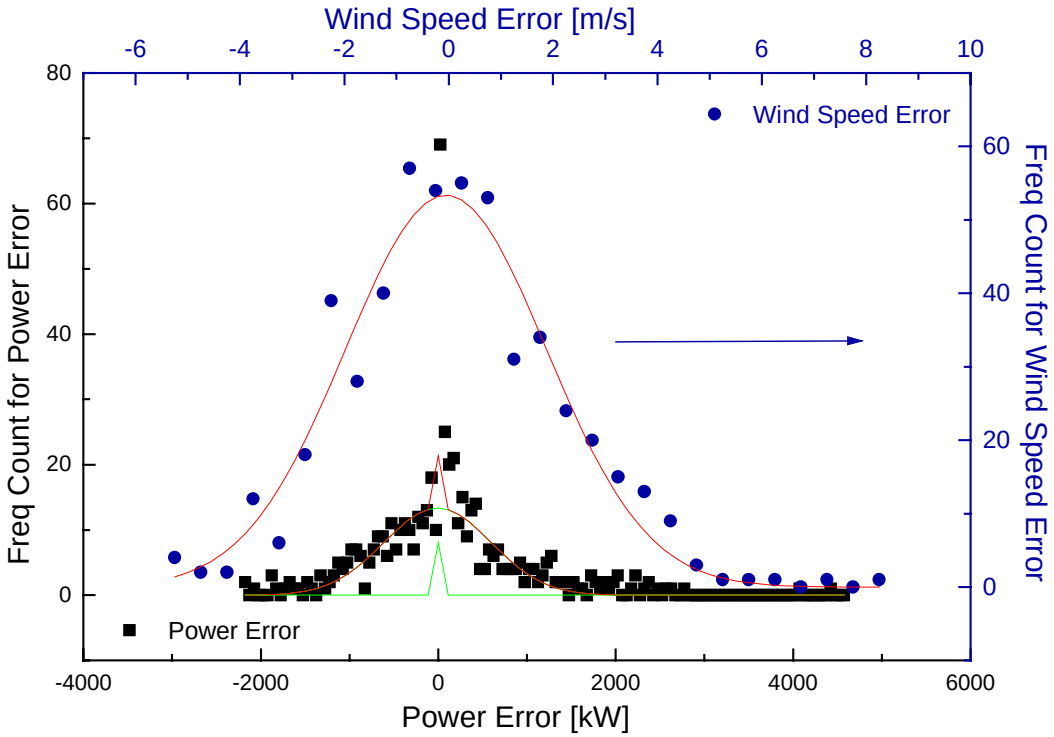

Figure 8: Frequency counts of the wind speed and power forecasting error for Nøjsomhedsodde wind farm, no MOS, 9 hours lead time. Peak power is 5175 kW. Please note the square at (0,70), which belongs to the narrower of the power peaks. The wind speed error is fitted with one Gaussian peak, the power error with two.

prediction can have quite significant error, without affecting the precision of the result.
Since the error is mainly Gaussian, the Kalman Filter is the optimal filter. However, this is strictly valid only for the use of a least squares criterion as quality function. We will see later how a different quality function affects the quality of the filter.

4.2 New Reference

The benchmark every effort in wind forecasting is measured against used to be the persistence model. This is a fairly good and hard to beat model for forecasting loads on the wind turbine out of the last developments of the on-site wind and for the development of corresponding control systems. However, these applications are only looking a few minutes into the future. With increased penetration of wind energy in the electrical grids up to quite large percentages, the utilities are craving forecasting models at least in the time scale of 8 hours, the time needed to start a thermal power plant from cold, or even up to 37 hours ahead for trading on the electricity markets.
Persistence is a good model for a short prediction horizon (it beats meteorological forecasts on a time scale of up to ~4 hours), but due to the processes in the atmosphere the deviation of this model gets higher for longer terms. For the prediction of long horizons it is intuitively clear that postulating the climatological mean wind gives a better forecast than a prediction from a random number, since a fully decorrelated time series is for prediction purposes nothing else than a random number with the same distribution as the predictand. Fully decorrelated means here that the autocorrelation function (see Appendix) of the time series has dropped to zero. Looking at Figure 10, Figure 11 and Figure 12, it can be seen that for most data shown here, the autocorrelation function drops to 0.2 within 48 hours. It can also be shown theoretically that the MSE for long term prediction by mean value is half the error for persistence [82].
Hence a new model should approach persistence for very short forecasting lengths and the mean for long times, while on the middle ground weighting between both models should occur. Another demand of a reference model is that it should not need any further input than the data set itself, and it should be much easier to use than any full blown forecasting model. These constraints are fulfilled in a proposed model where the weighting for different forecast lengths is determined by the correlation between the wind at present time and the wind the forecast length ahead. This reference model is usable as well for both wind and power. The model can be written

$$p_{t+k} = a_k p_t + (1 - a_k)\mu \qquad (12)$$

where p_k are measurements of wind or wind power at time step k; μ is the mean of the time series. The parameter a_k should be zero for $k=0$ and approach 1 with high k. Therefore it is reasonable to define

$$a_k = \frac{1}{N}\sum_{t=1}^{N-k} \hat{p}_t \hat{p}_{t+k} \Big/ \frac{1}{N}\sum_{t=1}^{N-k} \hat{p}_t^{\,2} \qquad (13)$$

with

$$\hat{p}_t = p_t - \mu\,. \qquad (14)$$

This actually corresponds to the value of a_k which minimises the RMS error for the reference model. This is also nothing else than the autocorrelation of the time series.

What exactly is long term? Figure 9 shows the skill score for the proposed reference model in comparison to the persistence model.

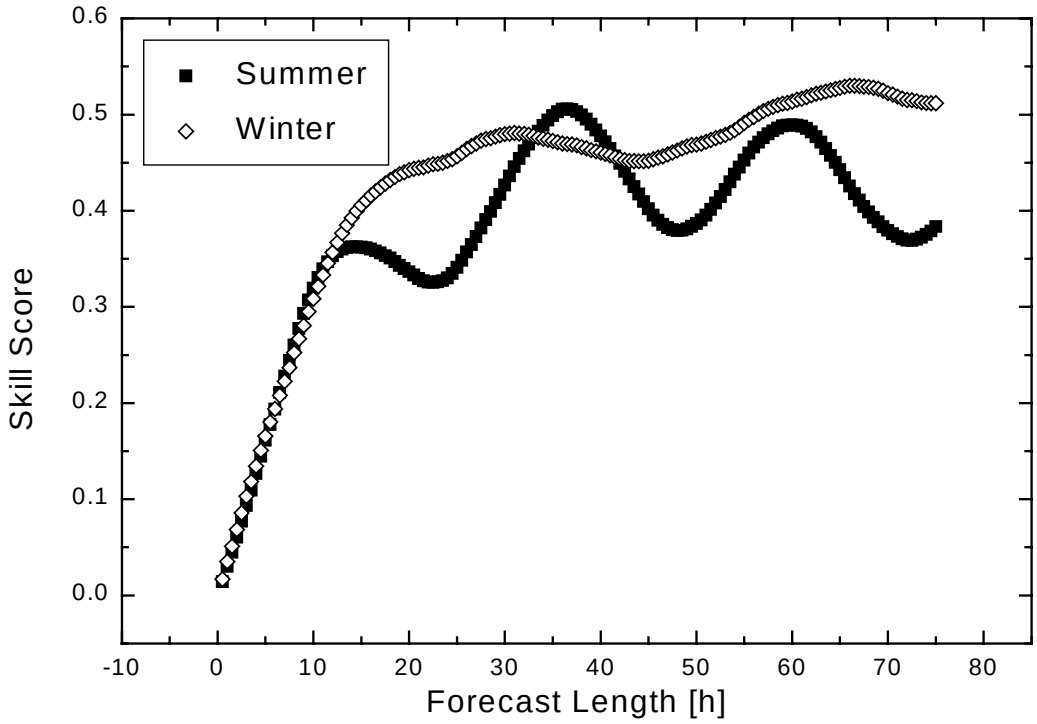

Figure 9: Skill Score of the proposed reference model compared to the persistence model. Summer and Winter refers to three months of production data each from the wind farm in Hollandsbjerg/Denmark.

In this figure the skill score (see Appendix) of the new reference model is compared to the persistence model using a data set of half hourly mean values of wind power from a wind farm located in Hollandsbjerg, Denmark. These measurements have been used to calculate the skill score as a function of the forecast length. Two datasets are considered, measurements from a summer and winter period. Each dataset contains 4380 measurements. After about 24 hours the skill score approaches 0.5, hence saying that the combined model has a MSE about half the one of persistence alone, and therefore the usefulness of the pure persistence *Ansatz* ends about here. On the other hand, a variation with a time scale of one day is easily discernible in the summer data, making the integration of the persistence model worthwhile also for longer periods.

Another feature found in Figure 9 is that the skill score rises even with the first step. This goes a long way to explain the typical 10% RMS error improvements found by many other authors using more or less advanced time series analysis techniques (see chapter 2.3). The nature of their algorithms takes some measure of the mean wind into account; either the mean of the last few data values in a windowing technique, or the weighted mean for many recent values using adaptive techniques. The improvement over persistence depends then on the number of recent data points used.

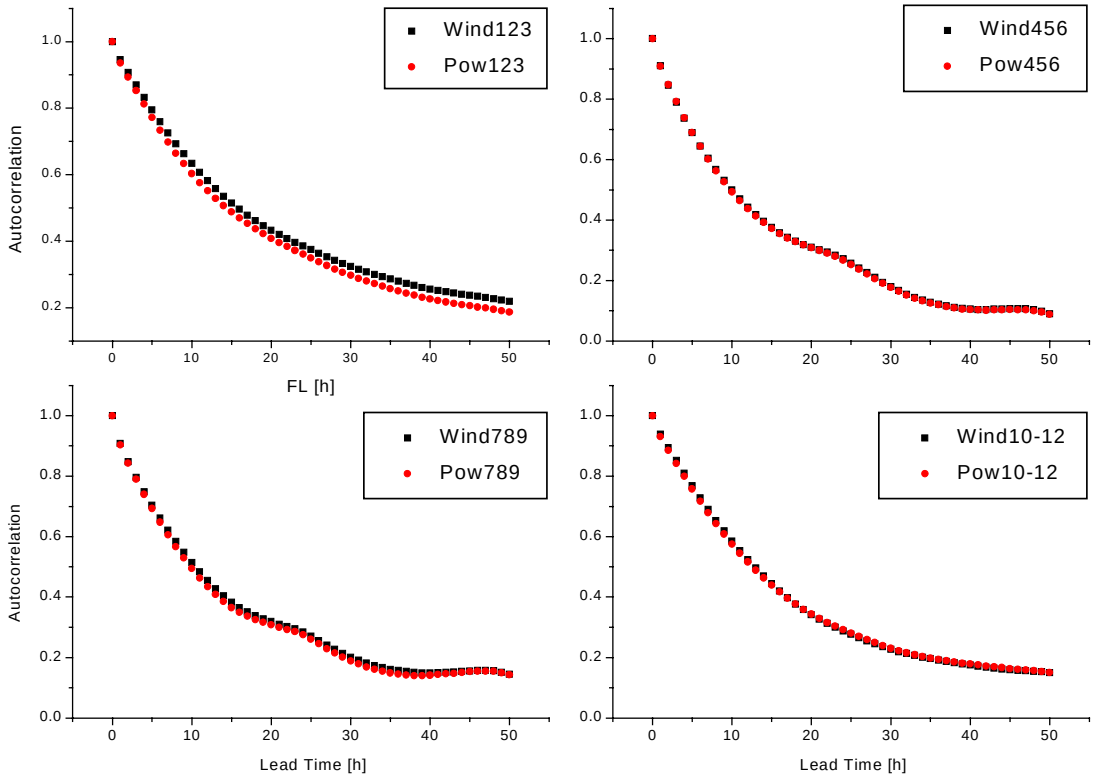

Figure 10: Autocorrelation of 40 years of data for the four seasons. The data is from 76m height a.g.l. of the meteorological mast at Risø. Wind123 refers to the wind from January, February and March. Pow refers to the same time series folded through a Vestas 225kW power curve.

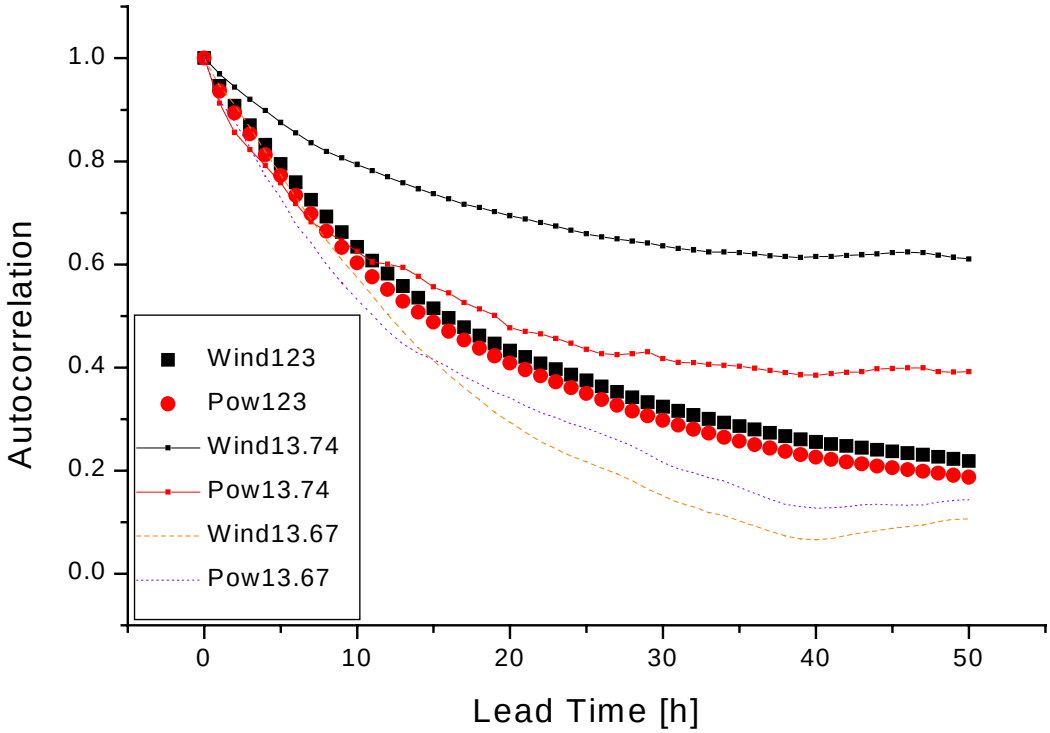

Figure 11: Some extreme cases of autocorrelations. 13.74 refers to winter 1974, 13.67 refers to winter 1967. Wind123 and Pow123 are identical to the graph shown in Figure 10.

It was tried to make the application of the New Reference Model easier for the user by omitting the autocorrelation step and parameterise the relative influence of persistence and mean by an exponential, which could be described as $p_{t+k} = a^k p_t + (1-a^k)\mu$. However, as can be seen from Figure 12, the autocorrelation for three sites varies significantly, from site to site and also from season to season; hence the "one parameter fits all"-approach was dropped.

The dependency on the season can also be seen in Figure 10. Here, the autocorrelation was calculated using 40 years of data from the meteorological mast at Risø. Another surprising result is that the autocorrelation varies also strongly from year to year: in Figure 11, the most extreme cases of autocorrelation functions of all single years are shown. The spread in just 7 years is quite wide. The *a* parameter for this case had been calculated, using this data and the European data described in chapter3.6, to be between 0.93 and 0.98 - good enough (like persistence) for the first few hours, but getting too far off for longer horizons (say, 24 hours).

In short, a new reference model is proposed, which is still easy to use, but overcomes some of the shortcomings of the persistence model for long forecasting times. It is possible to explain with this relatively simple model some of the improvements seen by many other authors trying to beat persistence on short horizons using time series analysis models.

4.3 Static MOS implementations

The easiest way to improve the results of a physical model is to apply a post-processor, trying to adjust the final output with a simple linear model of type *a*PhysicalModelOutput+b*. However, it might be desirable to have parameters elsewhere than at the back of the predicting algorithm. It will be shown here that adjusting the wind before calculating the power output is much better than any MOS performed solely on the power output.

The models used for comparison were[a]:

1) $Pow(a * u_{local}) * Eff + b$
2) $c * Pow(a * u_{local}) * Eff + b$
3) $c * Pow(a * u_{local}) * Eff + b * u_{local}^{3} + d$
4) $Pow(a * u_{local} + b) * Eff$
5) $a * Pow(u_{local}) * Eff + b$

Note: u_{local} refers to the pure HIRLAM/WAsP forecast which should be equivalent to the local wind at the site, the function *Pow()* is the power curve for the wind turbines used in the corresponding wind farm, and *Eff* refers to the corrections to the overall power generation due to wind shadowing effects in the farm, calculated with the PARK program.

Figure 13 is generally speaking typical for the errors encountered in the use of a weather model for short-term forecasting. In this figure, only the RMS error is shown, for simplicity reasons. The graph looks qualitatively identical for the MAE. The lines refer to different models: the uppermost line is

[a] $a * Pow(u_{local}) * Eff + b * u_{local}^{c} + d$ was tried, but the optimisation routine had problems with the numerical stability.

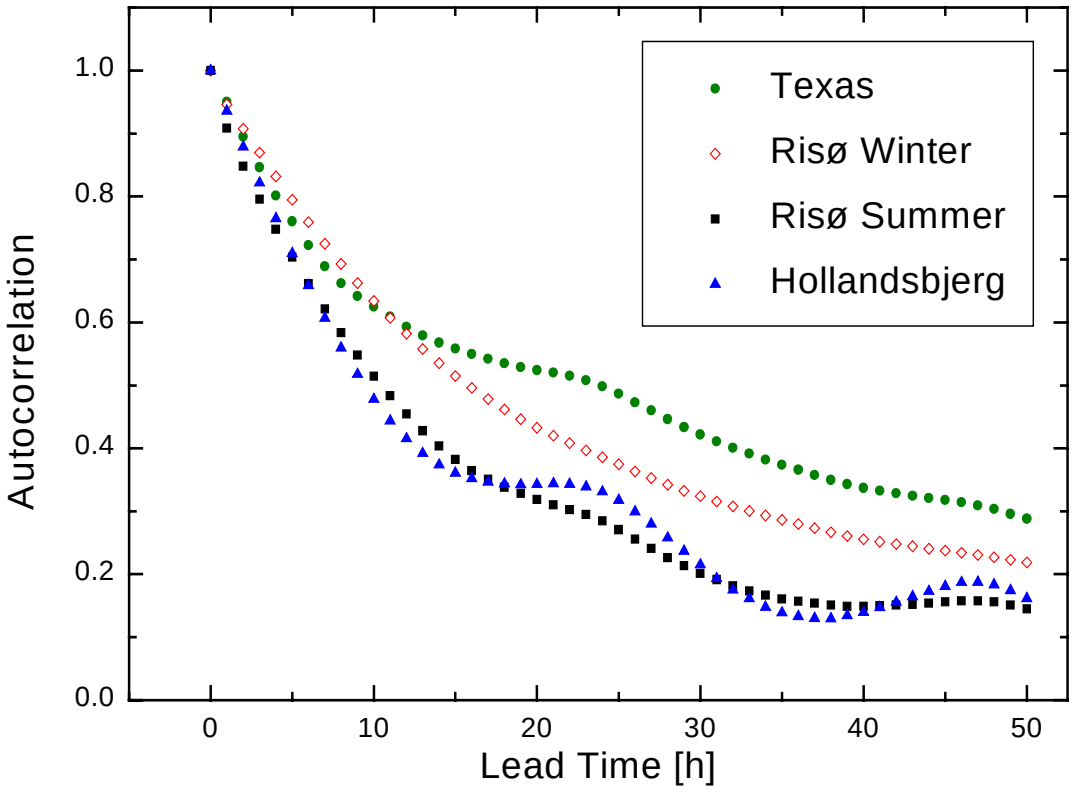

Figure 12: Autocorrelation of datasets derived from three sites. The Risø data comes from a meteorology mast on Sjælland/DK and is split into a summer and winter part, for each of which the autocorrelation was calculated separately. The Hollandsbjerg data is 4380 half-hourly values measured at a wind farm in Jutland/DK [83], while the Texas data stems from a prospective wind farm site; both series span more than one year. Note that Risø Winter is the only curve which can fitted satisfactorily with an exponential law.

persistence, the similar line below is the New Reference Model, both as explained previously. The scaling is chosen as RMSE in percent of the nameplate capacity, instead of the more statistically correct scaling to the mean of the time series, since for the utilities the total installed capacity in their supply area is much easier to keep updated than the mean production.

A few details should be pointed out:

- Mean refers to a forecast just by assuming that the wind x hours ahead will be like the long term climatological mean wind for the site. As can be seen, it outperforms pure persistence already at a forecast length of about 15 hours.
- Therefore, the New Reference Model was introduced. An improvement is visible over persistence on all forecast horizons.
- After about 4 hours the quality of the raw model output (marked HWP, full squares) is better than persistence even without any postprocessing. The quality of the New Reference Model is reached after 5 hours. The relatively small slope of the line is a sign of the poor quality of the assessment of the current state of the atmosphere by the NWP. However, calculating forward from this point onwards introduces hardly any more errors. This means that the data collection and the assessment of the current state of the atmosphere for the NWP is a weak point, while the mathematical models are quite good.

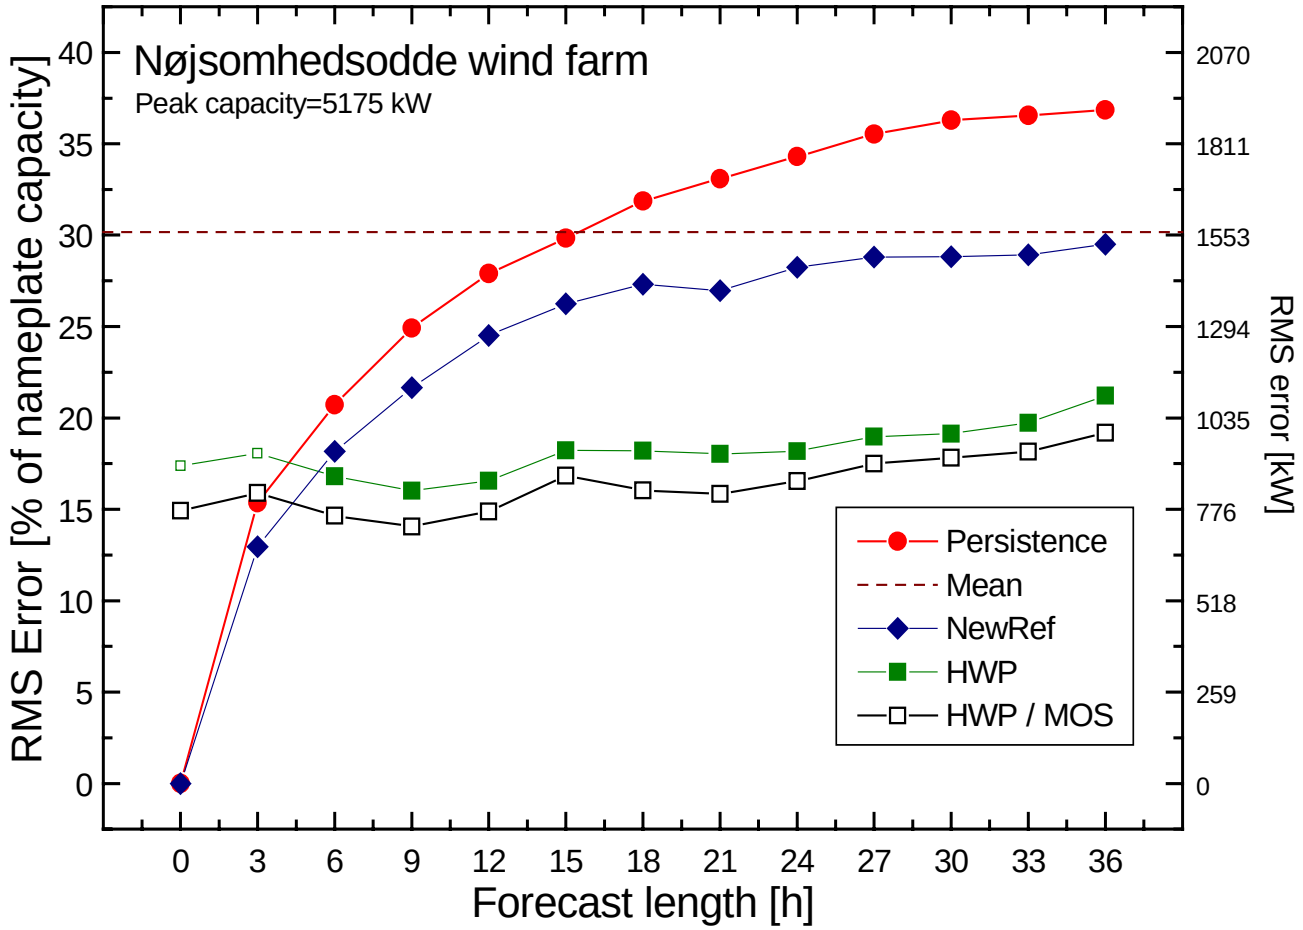

Figure 13: Root Mean Square (RMS) error for different forecast lengths and different prediction methods. The wind farm is Nøjsomhedsodde (noj) with an installed capacity of 5175 kW. NewRef refers to the New Reference Model. HWP/MOS refers to the HWP approach coupled with a MOS of type model 1.

- Obviously, a combination of persistence with the HWP approach would be helpful for the first few hours. This is planned for the next iteration of forecasting models from Risø. Here it has been tried, but since in the Risø operational set-up no online data is available, it could not be utilised.
- The first two points in the HWP line are fairly theoretical; due to the calculating time of HIRLAM (~4 hours) these cannot be used for practical applications and could be regarded as hindcasting.
- The improvement attained through use of a MOS (the line marked HWP/MOS, open squares) is quite pronounced.

The first try for a model would be an end-of-pipe-approach, trying to adjust the resulting power coming out of a black box modelling suite. This corresponds to model 5 in the previous list. In all cases the statistical parameters were optimised separately for every forecast horizon. Various approaches have been tried and are shown in Figure 14:

linear: an optimisation via the downhill simplex for minimum MAE;

corr.lin is the same, but the data set had been shifted by the mean error before the optimisation;

Kalman-20k has been an attempt to force the KF into saturation by running it 20000 steps through the same data set;

HWP-Pers is a linear combination of HWP and persistence;

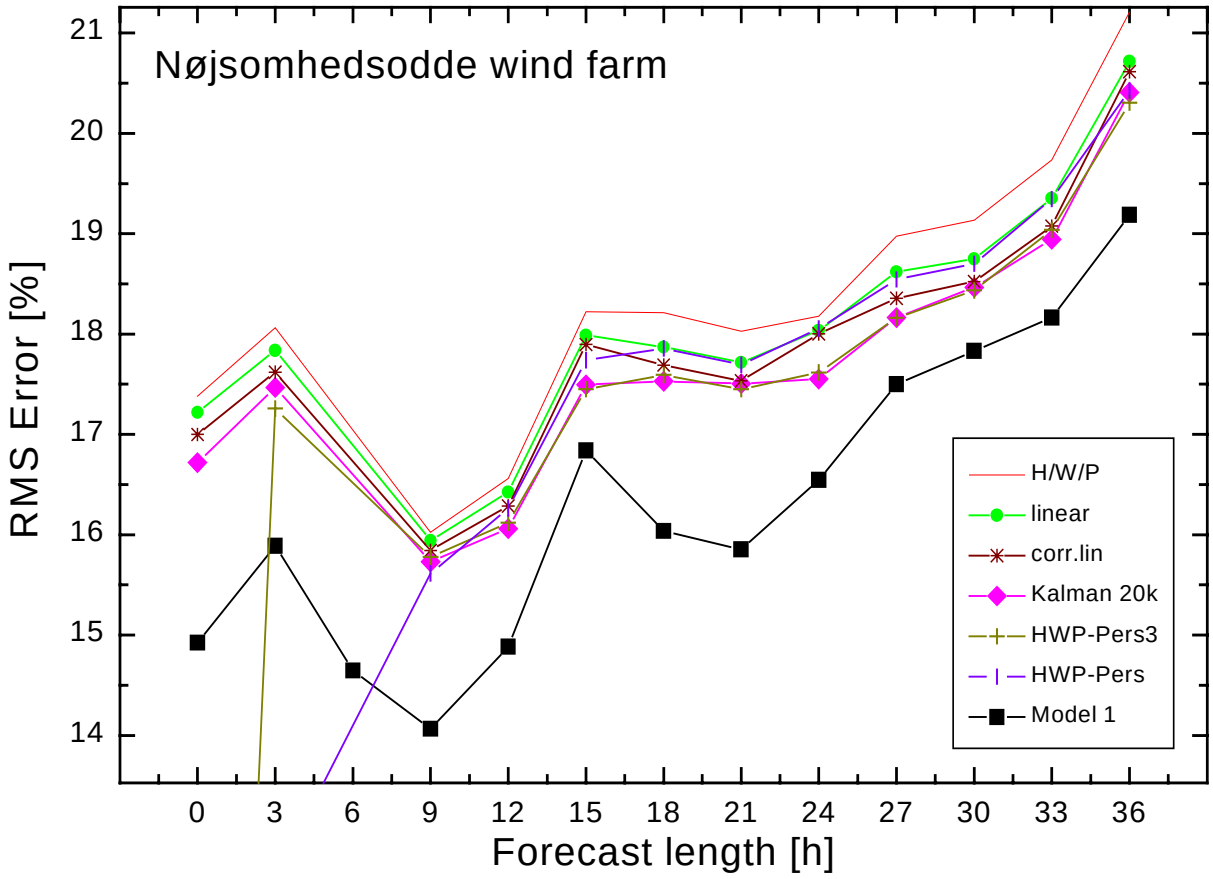

Figure 14: Root Mean Square (RMS) error for different methods of statistical improvements for the models. The wind farm is Nøjsomhedsodde (noj). Installed capacity is 5175 kW. New Reference refers to the New Reference Model. For the other models, see the description in the text.

and HWP-Pers3 is similar to HWP-Pers except for an additional constant. In Figure 14 only the RMS error is shown, for simplicity reasons. Again, the MAE is qualitatively similar, though smaller.

The main insights from this analysis were:

- The quality of the forecast and the improvements for the various methods do not vary greatly between each other. Since these were not optimised for a particular error type, the KF adjusted for the lowest RMS; hence some models give even worse results regarding the MAE than the raw model output (not shown).
- A model which combines the HIRLAM/WAsP/PARK results (filtered or not) with persistence including a weighting dependent on forecast length (marked HWP-pers) performs much better than HWP alone in the short range. However, for longer horizons no improvement is attained.
- The best result, both for MAE and RMS, is obtained by adjusting the HIRLAM/WAsP generated local wind before processing with the power curve and wind farm efficiency. This corresponds to model 1. This simple transformation yields visibly better results than all the statistical models applied to the end output alone. This also means that the HIRLAM generated wind is the main source of error in the chain of physical models.

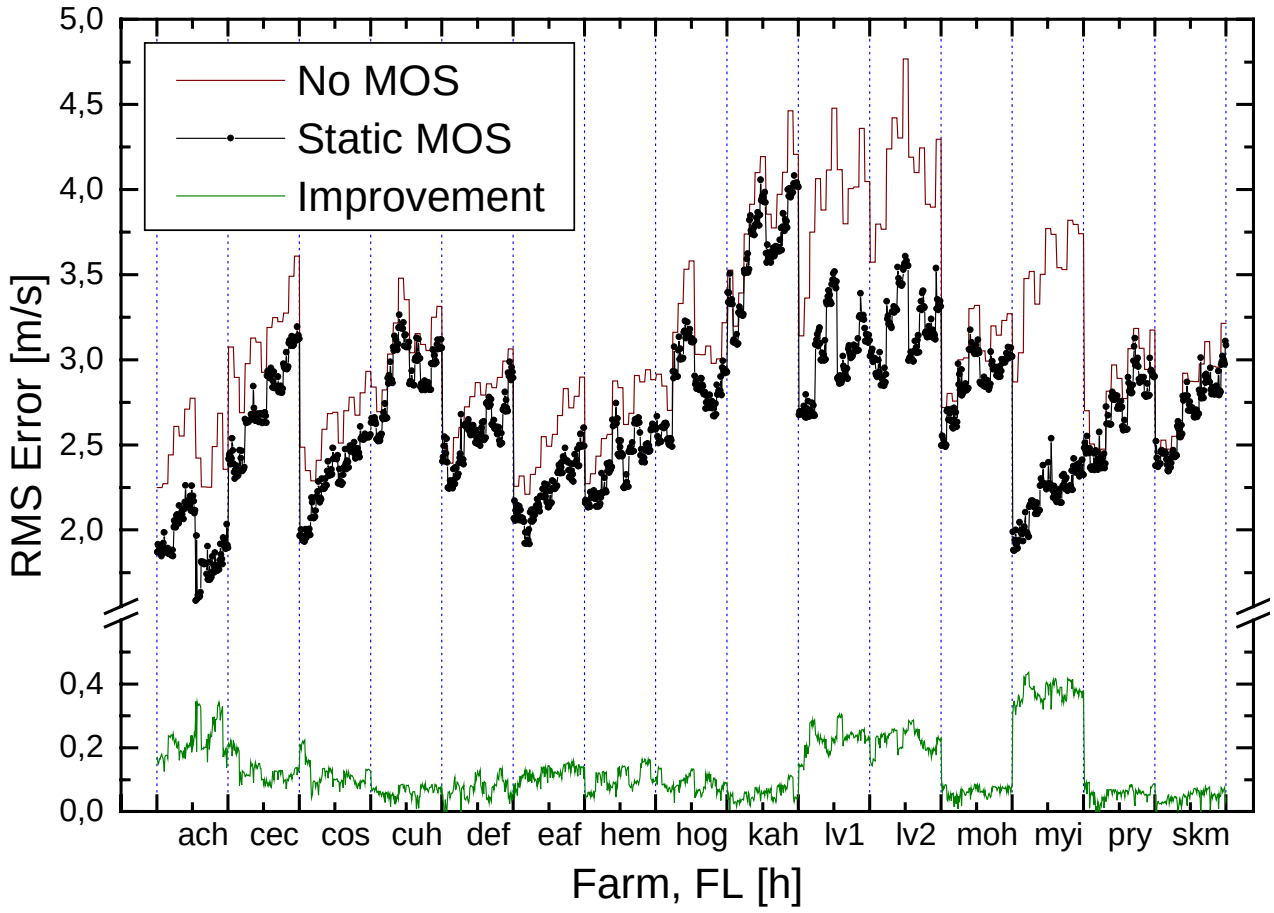

Figure 15: The relative improvement in wind speed RMS error for static MOS for 15 sites in England and Wales. Below the axis break is the relative improvement, given as 1-$RMSE_{MOS}/RMSE_{HWP}$. The x-axis lists the different stations, while within the vertical box for each station the different data points refer to the forecast lengths in steps of 3 hours, from 0 to 36 hours. This means that each of these boxes corresponds to the type of graph shown in Figure 13.

Since the performance of all these trials was significantly worse than model 1, they have not been pursued further.

Figure 15 shows the merits of a static MOS implementation for 15 farms in England and Wales. In most cases, the relative improvements in RMS error are fairly small, though for some farms the improvement can be much higher. This is most notably for myi, but also for lv1 and lv2; all of them are in Wales, and have quite a complex orography. This means that WAsP here is outside of its operational envelope. However, Bowen and Mortensen [20] have shown that the errors introduced by WAsP in the case of complex terrain pretty much scale linearly with the complexity. Hence, a simple linear MOS should be able to cover for them.

Another feature to be learned from Figure 15 (and Figure 17) is that the improvements do not scale with the forecast length. Then again, this could have been expected when thinking of the weak slope of the HWP graph in Figure 13: the biggest problem here was discovered to be the initial estimate, not the extrapolation into the future. Hence, the MOS improvement should be roughly the same all over, as it points to a systematic mismatch between the NWP and the local winds.

4.4 Kalman Filter results

Five statistical models have been applied to datasets from all 12 Danish wind farms using the following algorithm: firstly, the RMS and MAE were calculated for the plain physical models, then the best static parameters (*a*, *b*, and possibly *c*) were calculated with the downhill simplex method optimising for the different error types, and then the Extended Kalman Filter was run as MOS, here optimising the KF parameter via the downhill simplex method or a genetic algorithm in order to find the ideal "stiffness" of the filter. We will see to that in the next chapter in some more detail. The optimisation was done using both the RMS and the MAE as a quality function.

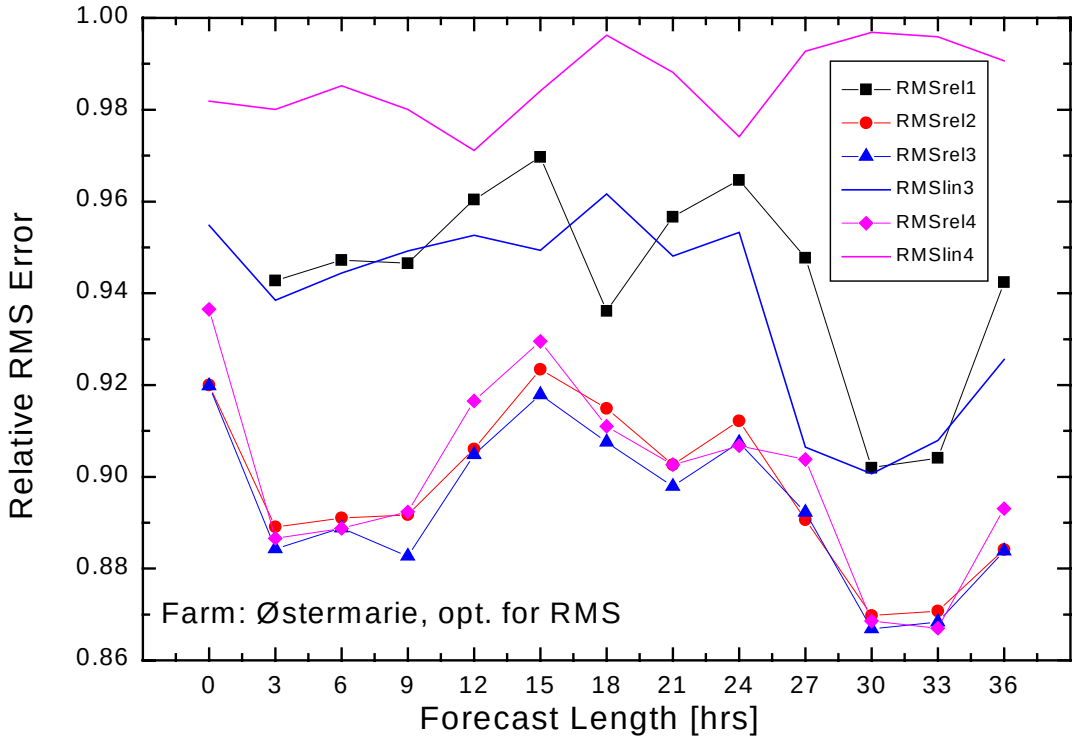

Figure 16: The RMS for models 1-4 for the site Østermarie (=oem), in the optimisation for RMS, relative to pure HWP modelling. Lines only refer to the model results with fixed parameters.

In Figure 16 the improvements for models 1-4 are shown for the Østermarie wind farm. The best attainable improvements seem not to depend much on the model, but rather on the quality of the pure physical forecast. Therefore, most models give similar results for the different forecast lengths. Models with a higher number of parameters (up to 4) were also tried, but did not perform better. Since the computational demands were much higher, it was chosen not to include these models in the current analysis. For this wind farm the improvement benefit of the KF is most pronounced - the same model with linear parameters gives only a smaller improvement, while most of the Kalman filtered MOS results give around 10% improvement. This can also be seen from the next plot: Figure 17 shows the results for model 4 when optimising for minimum RMS error. A few points typical of all models and error optimisations should be pointed out here:

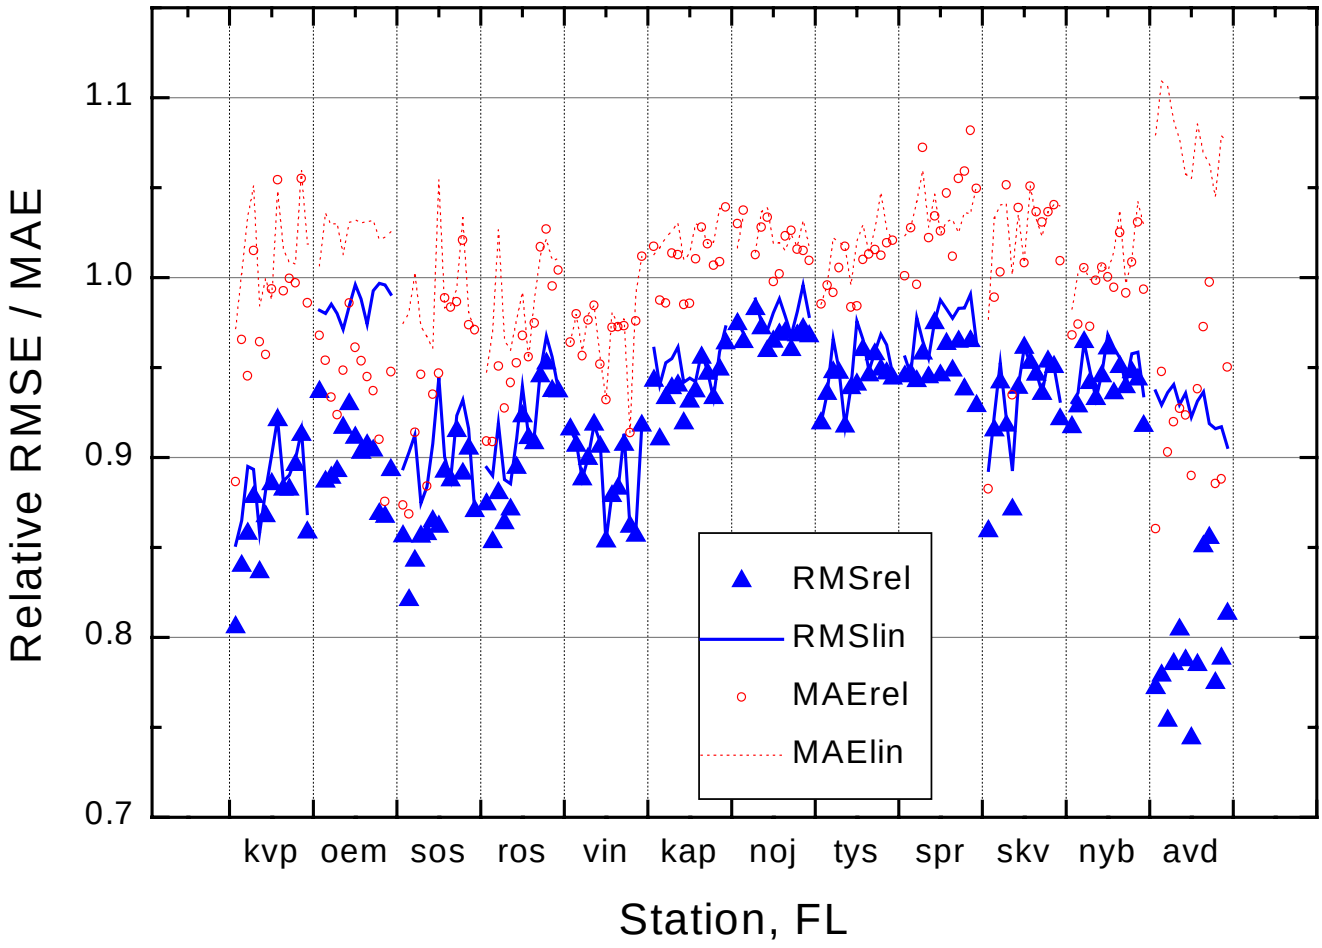

Figure 17: Relative improvements of model 4, optimised for minimum RMS error. Rel refers to the relative improvements of the Kalman filtering, while lin refers to the improvements due to linear optimisation of fixed model parameters. The open dots and dashed line refer to the corresponding MAE errors. Read the x-axis as in Figure 15.

- Optimising for best RMS error lets the MAE often get worse than with the unaided model. There seems to be a trade-off between the two errors - however, an optimisation for MAE usually also yields slight improvements for the RMS.
- Generally, the improvements achieved with statistical post-processing are in the range up to 20 %, which means that the underlying models already are quite good and exhibit mostly random errors, no systematic ones.
- Some stations are already predicted better than others by the physical models - and typically, the improvements the different models can achieve for these stations do not vary much. Here especially the stations kap, noj and tys are already forecasted rather well by the physical models.
- The advantage of the Kalman Filter is pronounced only for a few stations (this can be seen by comparison of the up triangles with the solid line for the RMSE), here mainly oem. Obviously, there is not much seasonal variation in the error of the physical models for the other stations.
- The case of avd is special, since here the power measurements were off during a few weeks. The optimal Kalman Filter could follow this jump in the relation between wind speed and power, though it gives no physical meaning. Taking out the bad data, avd behaves like the other farms. This also shows the importance of data cleaning in an operational forecasting model.

In Figure 18, we see the parameter trajectory for a complete run of model 1. This filter was found to present the optimal stiffness. Here, the seasonal variation of the parameters is, albeit small, clearly visible. In principle, these seasonal variations should be catered for by HIRLAM; reasons for the difference can be a different microclimate on the station itself than assumed by HIRLAM, or local effects such as heat flux induced stability variations, growth of vegetation in the seasonal cycle different from what is assumed in HIRLAM, or other unmodelled effects. The step structure of the variation can just about be seen. The allowed step width is determined by the stiffness of the filter and the scatter of the error. Also note that the overall change is very small. The KF tends to swing too much with higher flexibility. All it takes is one error value that is far out, and the KF wanders off in unphysical regions of the parameter space. Thereby, it sets all the forecasts to zero (at least in a model with a parameter inside the power curve, and we have seen that these give the best results). Once that happens, the feedback loop in the KF has no possibility of coming back on track. This is actually one of the biggest problems in the handling of the KF. One big disadvantage of using a stiff filter is that the time used to come into equilibrium is very large, hence the advantage of using a recursive method is nearly gone. Therefore, with the current error structure of the HWP approach, the KF is either so stiff, that the advantage over a static implementation is small, or follows the actual errors so fast, that it can go wild.

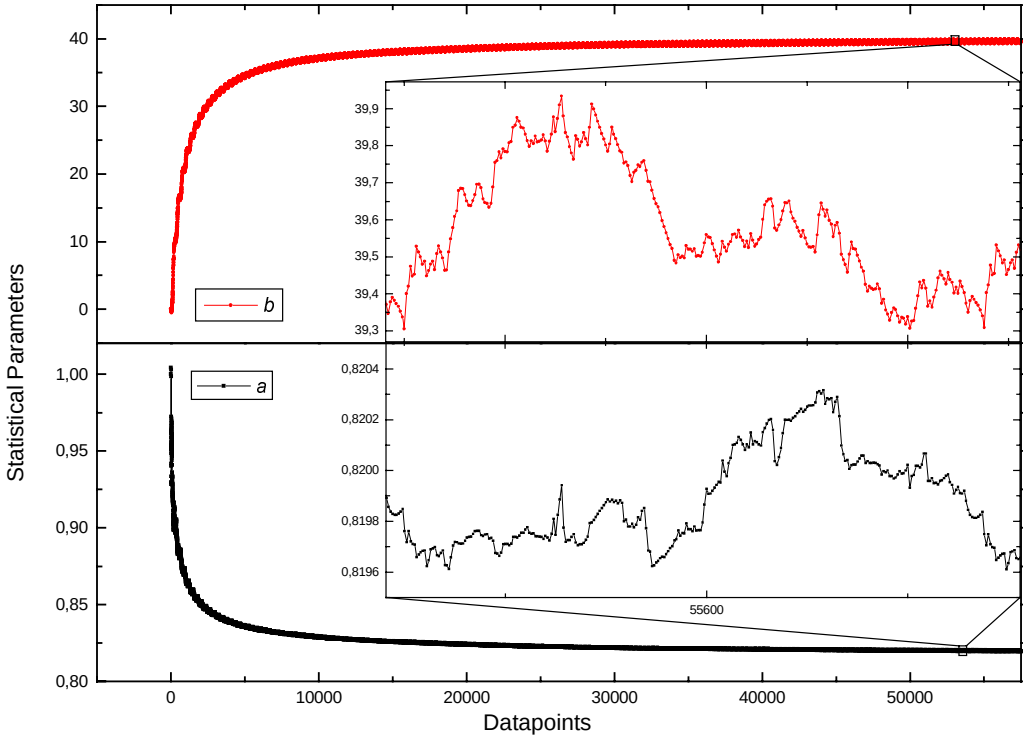

Figure 18: Parameter trace for a complete run. The inset is zooming in on one year's worth of data. Data from Vindeby, forecast length=24 hours, model 1.

4.5 Optimisation of the KF stiffness

As can be seen from Figure 18, it takes rather long for the KF to achieve balance. So since the stiffness of the filter is one of the parameters which can be used to fine-tune the filter, why not make the filter less stiff, so that it swings in faster? To analyse this, the best stiffness was optimised for using the downhill simplex and a genetic algorithm. Generally speaking, the genetic algorithm performed slightly

better, but took much more computing time. The set-up was to run the KF/EKF through all the data available, calculate the error function, run it again through all the data and repeat until the error function did not change any more. The general idea can be seen in Figure 18. Note that the typical set-up for the validation of statistical models of using one data set to calculate the parameters and another for the validation part is not applicable easily to recursive methods, since they keep learning during the validation dataset.

Figure 18 already points to the problem of the KF. The variation in parameters for the optimal stiffness is rather small, and therefore the potential for improvement over a static technique is rather limited. This has to do with the problem of the KF running wild. Large errors at one particular forecast will throw the filter off balance, and once a parameter within the power curve has run sufficiently far from reasonable values, the result will be plain zero for the power output. To show a complete mapping of the reasonable parameter space, in Figure 19 the minimum RMS error for the 24-hour forecast for Vindeby (vin) has been calculated depending on the parameters determining the KF stiffness. These are the diagonal elements of the measurement and system noise covariance matrices $\boldsymbol{V}$ and $\boldsymbol{W}$. For model 1 for low values of W_{22}, setting the power output irrecoverably to zero is what happens at the red plateau found in Figure 19. Model 1 has an additive constant to the power, which will become the mean value, while the multiplicative parameter adjusting the local wind has been thrown so much off course, that the result of all the HWP modelling is zero. Hence, this is the worst case. Another plateau found is the green plateau which is especially easy to see at $W_{22}=1$, but is also found in the other graphs. Here, the parameter variation is so small, that the cut-off criterion is reached already after a few runs. In other words, the filter does not move away from the initial guess for the parameters (which is total belief in the physical models).

The large variation in stiffness parameters necessary to achieve significant change is puzzling. No correlation with the properties of the time series could be found, and no correlation of the parameters with each other could be ascertained. The optimal parameters vary over many orders of magnitude for the same farm for different forecast lengths, and vary a lot from farm to farm. It was not possible to find one set of parameters for a model that would fit all farms. The probable reason for this is the low signal-to-noise ratio found in the deviations between the forecast and the measured wind. HIRLAM is already fairly good at also forecasting the seasonal variation. Mönnich [50] has shown that the atmospheric stability can be a factor determining the accuracy of the forecasts. However, the time scale of changes of stability is too short for the KF in relation to the noise. Of course, the stability should be modelled explicitly, though this was not possible with the available data.

Therefore, it has to be concluded that despite its scientific merits in reducing the forecasting error, as well as the good arguments for the choice of a recursive technique over a static implementation, the Kalman Filter has with the current structure of the HWP forecast errors hardly an advantage over simple static MOS implementations, but adds operational complexity. However, this should not rule out the use of this or another recursive technique for the future, when the HIRLAM modelling should be better, and therefore the signal-to-noise ratio should be improved.

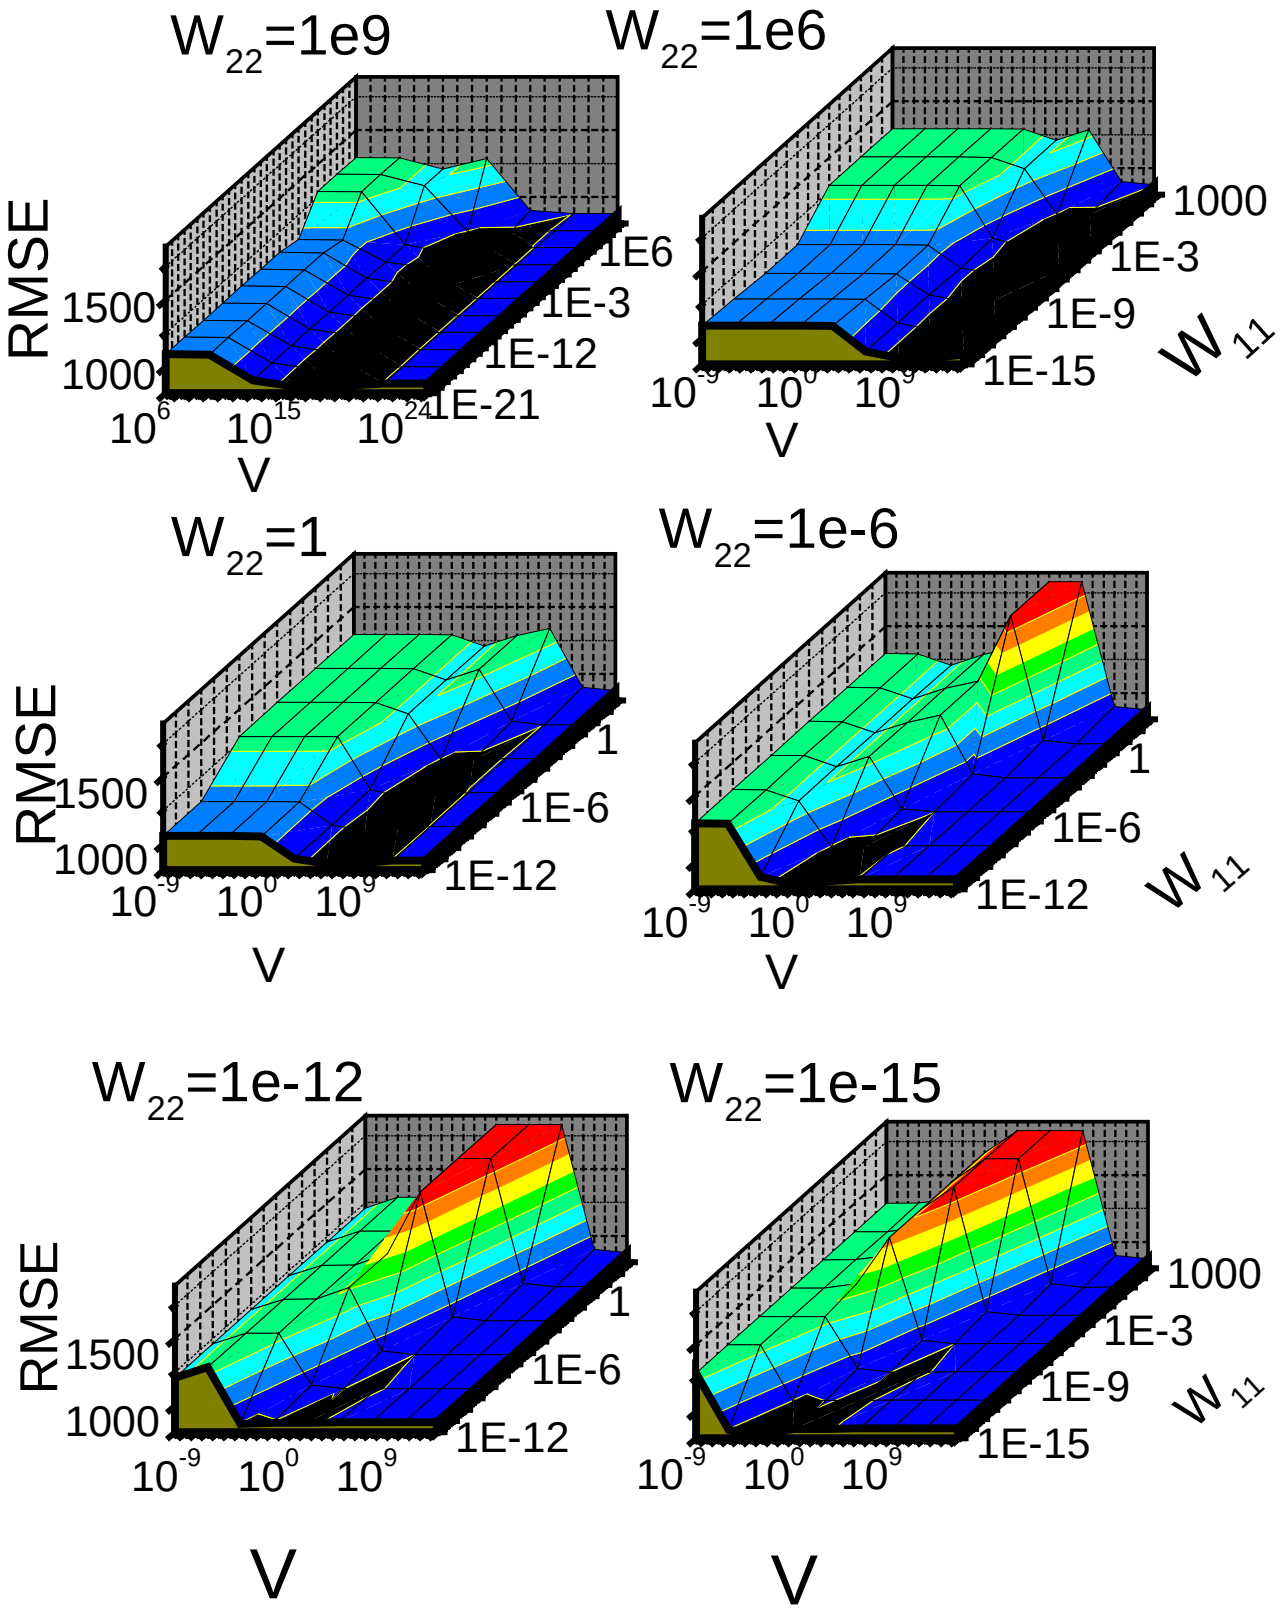

Figure 19: The RMS Error for different Kalman parameters. The farm is Vindeby, for 24 hours forecast horizon, model 1, optimisation for RMS error.

5. Benefits of good forecasting

5.1 Introduction

After the previous overview over forecasting methods, this chapter will deal with the potential benefits of forecasting. In the course of this chapter, the National Grid Model should be made familiar for the reader, and some answers regarding the sensitivity of the model to variations in turbine type, plant mix and load magnitude should be given here, before the NGM is subsequently used to analyse the whole of Europe. Additionally, an impression of the benefits of perfect forecasting over persistence and the penetrations, at which these appear, are given.

On an intuitive level, the benefits for the utility dispatcher appear clear. Since the load also cannot be predicted with perfect accuracy, all utilities have the possibility to adjust their production on very short notice. This is also necessary for the case of the unexpected loss of a major power plant. Typically, this load-following elbow space is achieved by spinning reserve or warm back-up stations. Also pump-storage plant has a good potential for fast variation. For the fairly predictable changes of the load that occur on a time scale of hours, another, slower mechanism has to come in: scheduling. The system operator can tell a power plant to start up or shut down so that some additional capacity is at the command of the operator. The typical time scales here are between 20 min. for gas turbines and 8 hours (or perhaps more) for large coal or oil plant. The order in which these plant are scheduled is a merit order, which means that the power stations with the lowest cost per kWh_{el} are the ones starting first and stopping last. Optimal system commitment is achieved if at any given time the minimum number of power plant needed to serve the load is running. This should include a safety margin to allow for the sudden drop of a power source (due to technical failure in a power plant or the transmission line leading away from the power station, or due to a technical failure of the interconnector to another network[a]). Note that in a real-life situation, the penalty for overcommiting might not be the same as the penalty for undercommiting. Typically, the price will be slightly asymmetric for buying or selling electricity. Note, too, that in a system as large and as well interconnected as the European grid, some extra power should be available at all times, even though it may have a price[b].

One simple conclusion is that as long as the installed capacity of the intermittent resource leads to a uncertainty in its output that is in the same order of magnitude as the uncertainty in the load, the utility should be able to make up for the added uncertainty with the existing tools. Conceptually, this is treating wind power as negative load while not taking it into account at the scheduling stage. However, with higher installed capacity this approach does not work well any more.

There are two possible reactions to the integration of intermittent power sources: one could totally disregard any firm capacity and schedule ahead as if no power from that source could be expected, or

[a] The biggest "power station" in the CEGB grid *eg* is the interconnector to France, with 2x1 GW capacity.

[b] One of the highest prices on NordPool since the liberalisation was 90.13 €c/MWh on December 12, 1999, during a very cold snap across the Nordic region (Source: Financial Times, 18 Jan 2000). Even more, the price in England and Wales regularly reaches during winter peaks £300/MWh.

one could try some forecasting and would have to deal only with the remaining error. If typical gradients are known [*eg* 84], then to be on the safe side a line from the actual production can be drawn with the maximum downslope, until it reaches zero production. This would lead to an inclusion of wind power for only a few hours ahead. However, even a simple analysis shows that on average the operators best choice is to use a forecasting model close to persistence in the absence of any further information. Nevertheless, it can happen that the dispatcher does not even know the actual power feed; in this case, the pure HWP modelling would be the forecasting model of choice.

Even though the case for forecasting (on which level whatsoever) is an easy one, there are not many analyses that have looked in detail into the benefits of forecasting for a utility. Partly this lack of analyses stems from the fact that a lot of data input and a proper time step model are needed to be able to draw valid conclusions.

Milligan *et al* [85] used the Elfin model to assess the financial benefits of good forecasting, taking into account the load time series, a wind time series, the distribution of power plants for different utilities, and the forced outage probabilities of the normal plant mix. Even though his method of simulating the forecast error was not very close to reality, some general conclusions could be drawn. When varying the simulated forecast error for three different utilities, zero forecast error always came out advantageously. The relative merit of over- respectively underpredicting varied between the two utilities analysed in detail: while underpredicting was cheaper for one utility, the opposite held true for the other. The cost penalty in dependency of the forecast error was dependent very much on the structure of the plant mix and the power exchange contracts. Generally speaking, a utility with a relatively large percentage of slow-start units is expected to benefit more from accuracy gains.

Hutting and Cleijne [86] analysed the proposed structure of the Dutch electricity exchange, and found that 1500 MW of offshore wind power could achieve an average price of 3.5 €c/kWh, when coupled with back-up conventional plant. This assumes that "*75% of the output can be predicted well enough for the market*". Perfect prediction would raise the price to 4 €c/kWh. However, building 6000 MW of wind power would decrease the price to 2.9 €c/kWh. Reducing the specific power of the rotor from 500 to 300 W/m^2 would decrease the overall power output, but increase the capacity factor, thereby increasing the predictability and therefore enhancing the value by an extra 0.05 €c/kWh. This would actually improve the price performance ratio by about 10%, just by installing larger blades on the turbines. Spreading out the wind farms along the coast would increase the reliability of the generation and therefore lead to another 0.15 €c/kWh.

Nielsen *et al* [87] assessed the value for Danish wind power on the NordPool electricity exchange to be 2.4 €c/kWh in a year with normal precipitation. This would be reduced by 0.13-0.27 €c/kWh due to insufficient predictions.

Sørensen and Meibom [88] found the penalty due to bad prediction of wind power to be 12% of the average price obtained on NordPool, which is essentially the same as the result by Nielsen *et al*.

Schwartz and Brower [89] interviewed schedulers, research planners, dispatcher and energy planners at seven US utilities and asked for their needs in a wind energy forecast. Among the most needed was a day-ahead forecast, to be given in the morning for the unit commitment schedule and energy trading for the following day. Hourly forecasts, expressed in likely MW and with error bars, were another wish.

However, one important result was that if good tools were available, operators in utilities with enough penetration would use these tools. This is also our experience with operators from Danish utilities.

5.2 NGM – Sensitivity analysis

Even though the National Grid Model is a quite old model, a thorough analysis of the sensitivity to the type of turbine has never been carried out.

For this analysis, measured wind speed from prospective sites in Iowa were transformed into wind power outputs using power curves supplied by the manufacturer. The turbines used were a Nordtank rated to 150 kW, the Vestas V27 with 225 kW, a 400 kW turbine, a Micon M750 (750 kW), another Nordtank turbine with 1500 kW peak and the Vestas V66 with 1650 kW rating. For some configurations, also a NEG Micon turbine with 1000 kW was used. These are also the turbines that appear in the European wind turbine distribution used for the European average wind in chapters 3.6 and 7.

These different forecast files were then used as input for the NGM. In Figure 20, the results for 1995 and 1996 for Iowa are shown as full symbols. The forecasting method here is persistence. To assess the influence of the load and the wind power time series, these two were interchanged - hence the 1995 load curve was run with the 1996 wind power time series and vice versa. Note that the plant mix was not changed for comparability reasons, even though one new plant entered service in 1995, an open cycle gas turbine with 15.2 MW. For comparison, there are 122 open cycle gas turbines in the whole grid with a cumulative generating capacity of 1274 MW. Hence, the contribution of the new turbine is negligible.

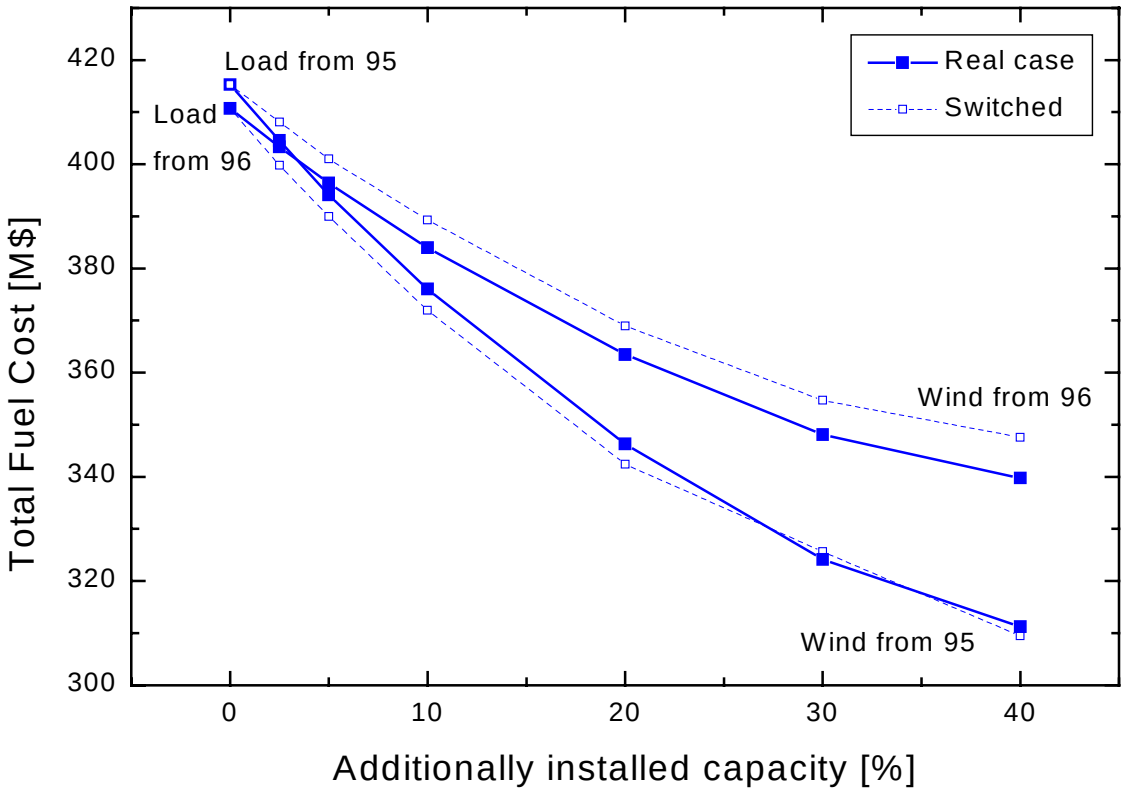

Figure 20: NGM results for Iowa, with varying loads and winds. Loads from 1995 and 1996 were used interchangeably with the respective wind time series.

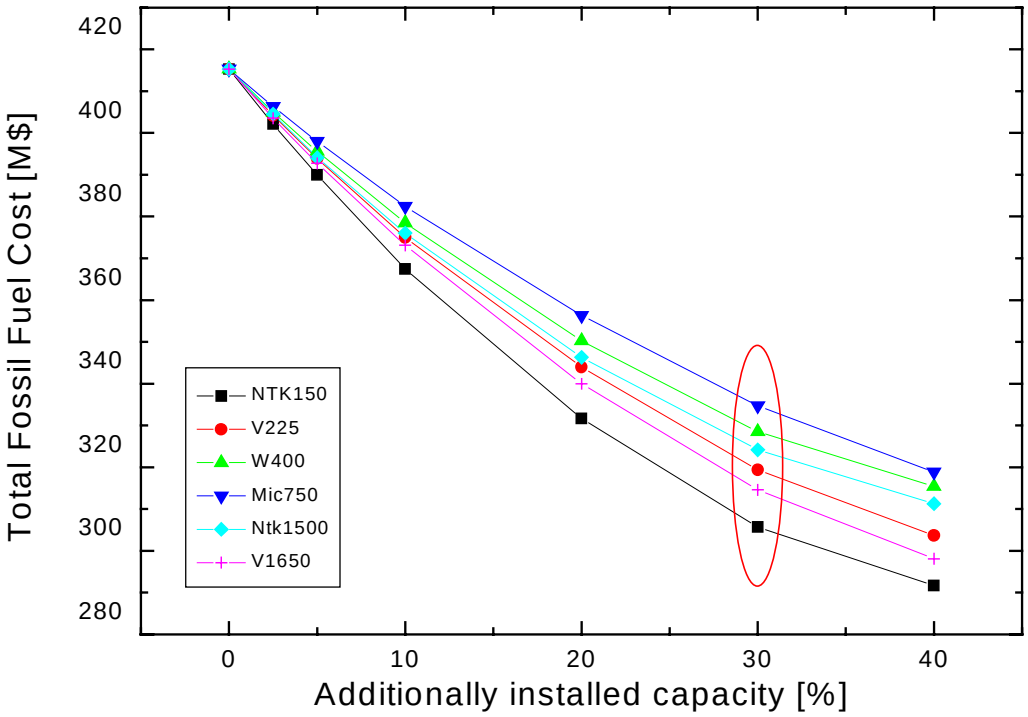

Figure 21: Total fuel cost depending on different turbines. The circle refers to the points used for the analysis in the following figures.

The influence of wind and load time series is shown with open symbols. As one might expect, the shape of the curve is determined by the wind time series, while the magnitude is mainly determined by the magnitude of the load. The most interesting feature here is not the apparent parallelism of the curves with wind from 1996, but the overlap for high penetrations in 1995. This stems from the relative distribution of the wind and load time series during the year. Interestingly enough, if one uses a load multiplication factor determined from the cumulative load and multiplies the whole time series with that factor, the result for 0 penetration (and all higher ones) is not identical with the load curve from the other year. Obviously, there are other factors involved controlling the total fuel cost over a year than the total demand over that year. However, as before the shape of the curve is governed by the wind data file while the magnitude is governed by the load. The difference is explained most simply with the wind conditions in the respective years: in 1995, wind generation for a 1500 kW turbine would have been 3.914 GWh, which translates to 2609 full load hours. The figures for 1996 are only 2.722 GWh and 1815 hours, respectively. This difference is because many more measurements from 1996 were missing than for 1995. The wind time series is composed by taking the mean of all the wind sites available, in this case just two sites. Missing data in only one site is hence straightened out with the value from the other farm. However, if both values are missing this is treated as a turbine outage, hence the value is set to 0. This happened 116 times in 1995, but 726 times in 1996, hence nearly 10% of all data was missing and set to zero. Unfortunately, most of these 10% were in the winter, when strong winds prevailed. Hence, the aggregated output for 1996 was significantly lower. Even when using only the available points for a comparison, 1995 has the stronger generation with a capacity factor of 29% versus 24% in 1996. This is because the stronger winter winds are lacking in the 1996 statistics. However, this variation can happen from year to year. See *eg* the difference between the average European production in 1986 and 1987 in Figure 49. Keep in mind that the European production is averaged over a large area,

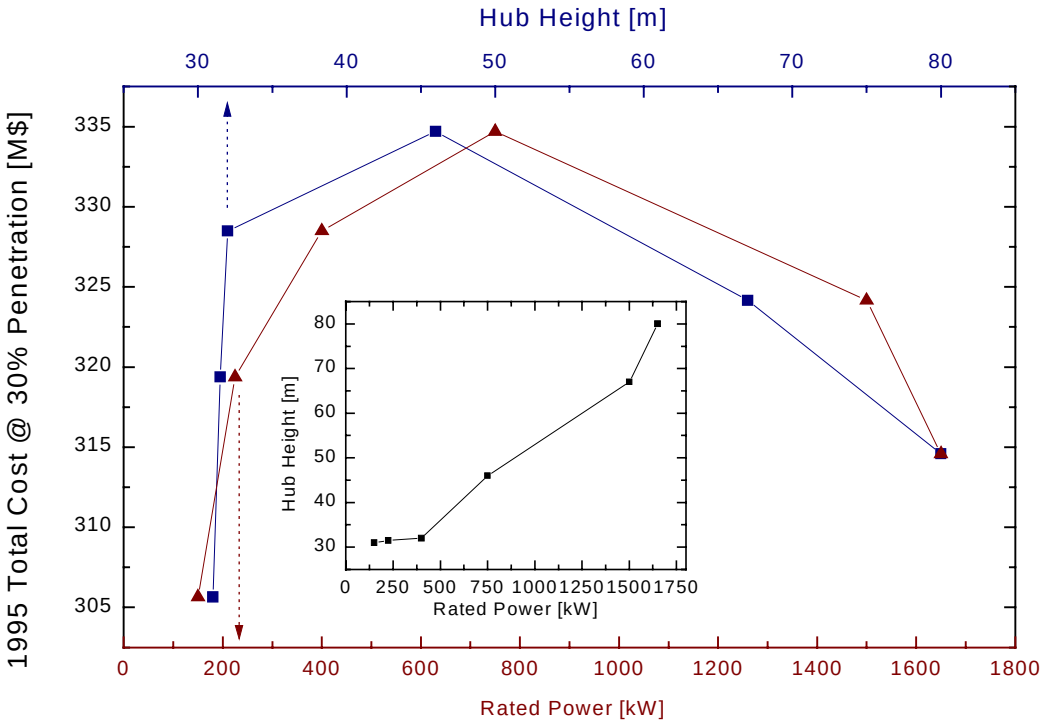

Figure 22: Total fuel cost for different turbines. The inset shows the dependency of height and rated power (a higher rated power, the turbine needs to be larger).

and therefore should be less variable than the variation for just two sites in the same climatological region here.

All the previous analyses were done using a Nordtank 1500kW power curve. So, what happens if we use other power curves? This was evaluated for 1995.

Note that there is no clear trend with size: the lowest fossil fuel savings occur with the middle sized turbine, while the best savings are for a 150 kW and a 1650 kW turbine, respectively. This is shown in more detail in Figure 22. Here the values in the circle in Figure 21 are shown over the rated power of the turbine, or over the hub height. We see that the biggest savings (lowest overall costs) occur for the smallest and the biggest turbines. The two curves are similar, since the correlation shown in the inset is nearly linear - the power output of a turbine is related to the area swept by the rotor, hence a taller turbine allows for bigger rotors and hence a higher power output.

This behaviour needs an explanation. For a start, the power curve used here is shown in Figure 23. Actually, as can be seen from Figure 22, the Nordtank 150kW is nearly the same size as the Vestas 225kW turbine. In other words, the generator in the Nordtank turbine is too small for the blades, and hence will have a higher capacity factor. However, in the case of the 150 kW machine, another bug comes into play: the power curve levels out at 159.7 kW, hence the turbine should be called a 160 kW machine. The manufacturer has decided otherwise, therefore it performs so outstandingly in this comparison. Note however, that the Vestas V27 has a rating of 225 kW and also levels out at 225 kW, and still has the edge on slightly bigger machines. See also Mönnich [50] on the comparison between a theoretical power curve and a measured one.

Then again, this only explains the first half of Figure 22. Another mechanism is responsible for the higher savings attainable with bigger machines: the wind is higher at higher altitude. As can be seen in

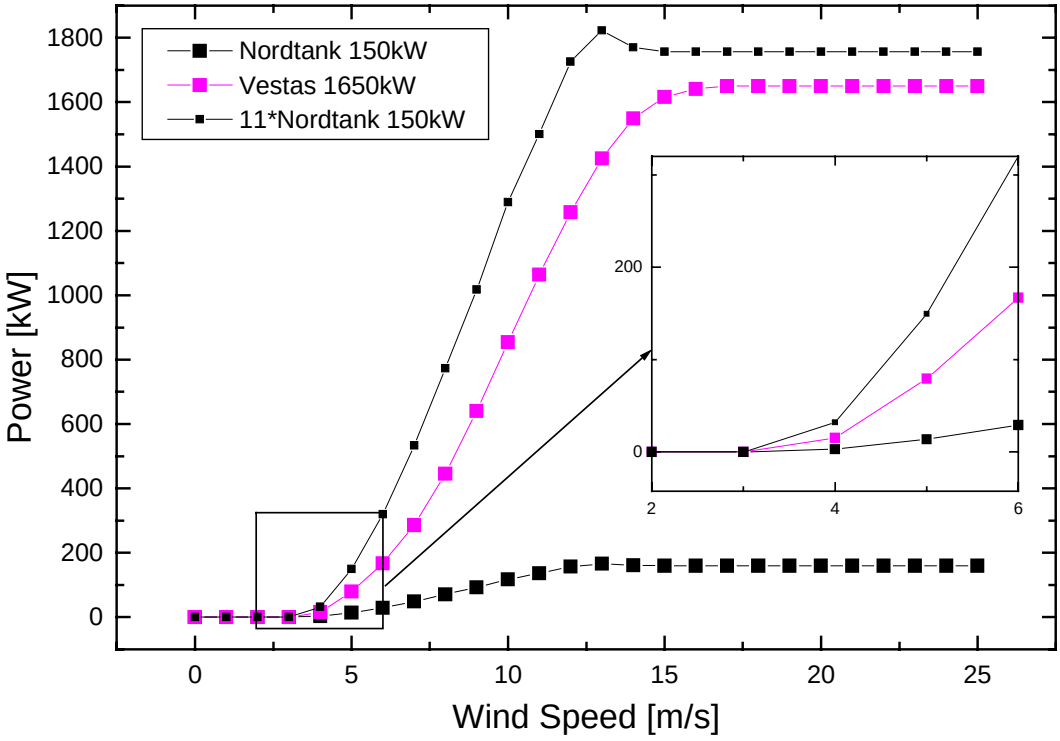

Figure 23: Power Curve of a Nordtank 150kW, and the comparison between a Vestas 1650 kW and an array of 11, 150 kW machines.

the inset of Figure 24, the higher and bigger machine taps a greater resource. The main graph shows the simple explanation for the different total fuel costs from Figure 22: the amount of energy generated from the same wind is different, due to the two effects of higher specific rotor area per power rating for smaller machines and higher wind resource for big turbines.

This leads to the conclusion, that the potential savings of fossil fuel cost scale with the accepted wind energy. The accepted wind energy is the generated wind energy minus the amount of energy that has to be discarded. This happens when at some hours of the year the base load generation plus a fraction of the available wind power suffices to cover the demand. The remainder has to be discarded. The difference is only visible for additional capacities nearing 20% of the currently installed conventional capacity. Then, the differences between the production of different turbines begin to appear in the fossil fuel savings. The curve based on accepted wind energy instead remains clear.

Another proof for the scaling with accepted wind energy can be found in Figure 26. Here the use of three different turbines (a 225kW model, a 1000 kW and a 1650 kW) does not make any difference, nor do three different wind time series, or different forecasting methods. Neither does the scaling of the time series with a factor of 1.4, which was done to get a rough impression of what the wind would be in 60m a.g.l., the hub height of newer turbines in the megawatt class. This actually corresponds to applying the logarithmic wind profile over a uniform roughness of 0.12m.

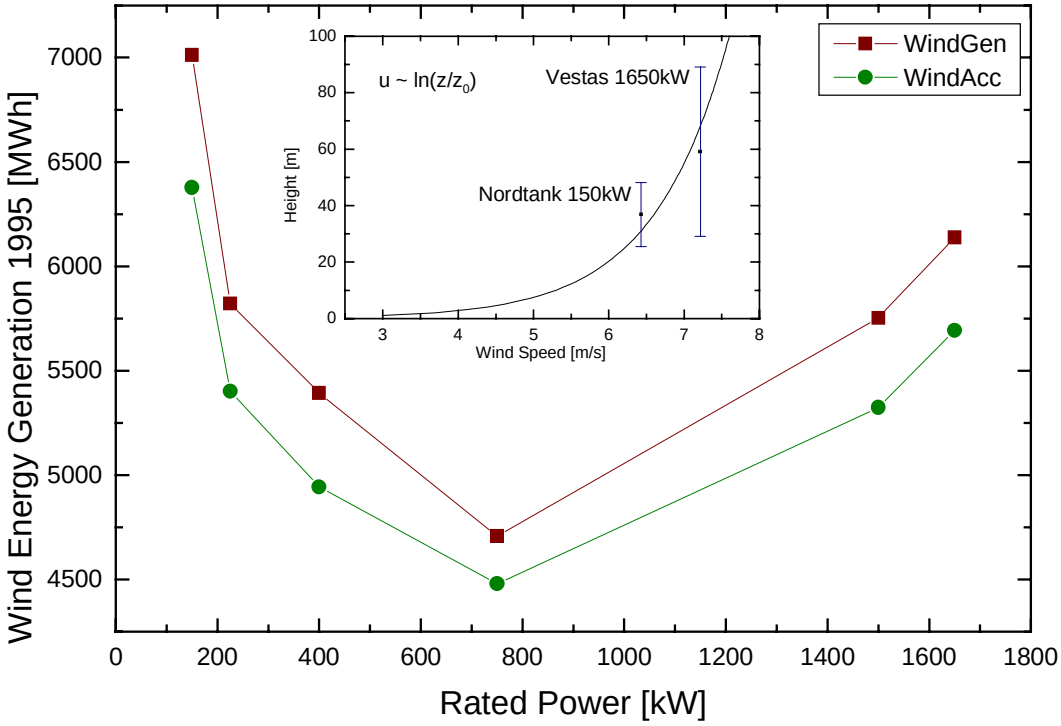

Figure 24: Electrical energy generated by wind in 1995. WindGen refers to generated wind energy, while WindAcc denotes the amount of wind energy that was accepted into the grid. The inset shows the logarithmical height dependency of the wind, together with the relative sizes of the Nordtank 150 kW machine and the Vestas 1650 kW.

An interesting result in this case is that the difference between perfect forecasting and persistence forecasting is quite low. An explanation for this behaviour is that the way of forecasting influences the amount of discarded energy (see *eg* Figure 34). Perfect forecasting allows a higher fraction of the generated wind energy into the grid, whereupon the attainable savings correlate with the amount admitted to the grid.

The plant mix has a strong influence on the result (Figure 25). The largest difference between the plant mix of 1994 and the one of 1988 is that in 1994 the 'rush to gas' already had happened. Therefore, the flexibility of the receiving grid was augmented significantly, and wind energy could be integrated much more easily. However, this explanation does not help as to why the recasting of gas turbines as steam plants has lower total fossil fuel costs than with them. The probable explanation is that the gas turbine capacity was recast to state-of-the-art coal plant, with accordingly rather low cost/kWh$_{el}$. The fuel costs for gas turbines are much higher than for coal fired power stations.[a] This effect could easily offset the higher spinning reserve requirement in a grid with less flexibility.

[a] Only the relatively low initial investment for building the plant makes gas turbines so attractive. However, the investment cost was not part of the analysis.

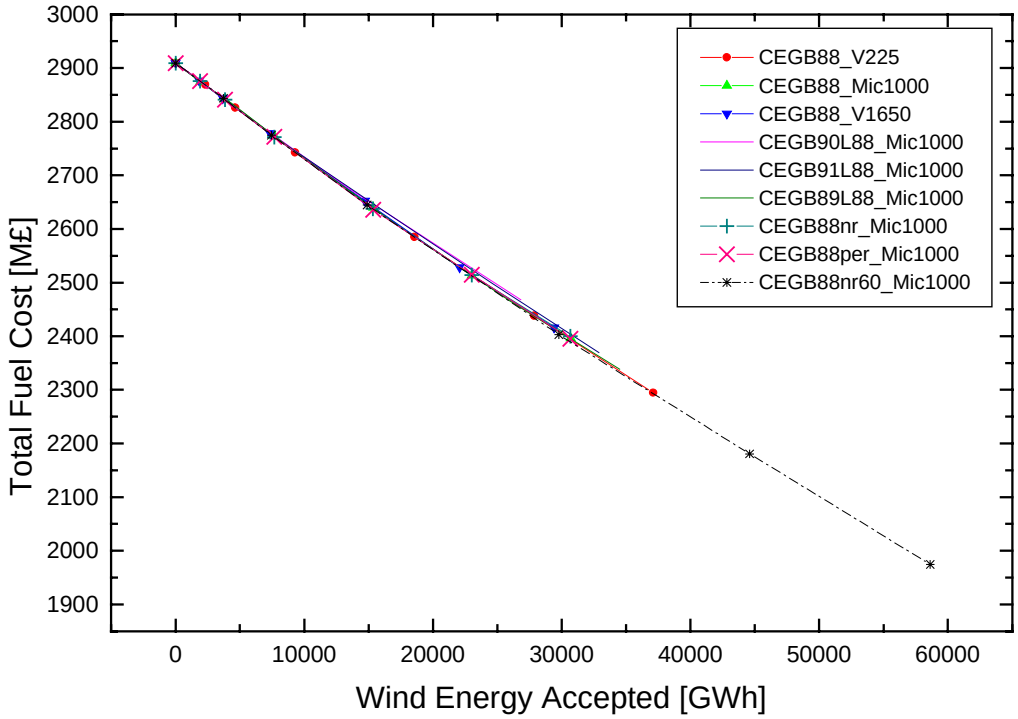

Figure 26: Total fuel cost for the England and Wales grid for the 1988/89 load profile. V225 refers to a Vestas 225kW turbine, Mic1000 to a Micon 1MW and V1650 to the Vestas V66 with 1650kW peak generating capacity. 90L88 refers to the wind year 1990, combined with the load year 1988. 88nr refers to the New Reference Model as forecasting method, while per stands for perfect forecasting.

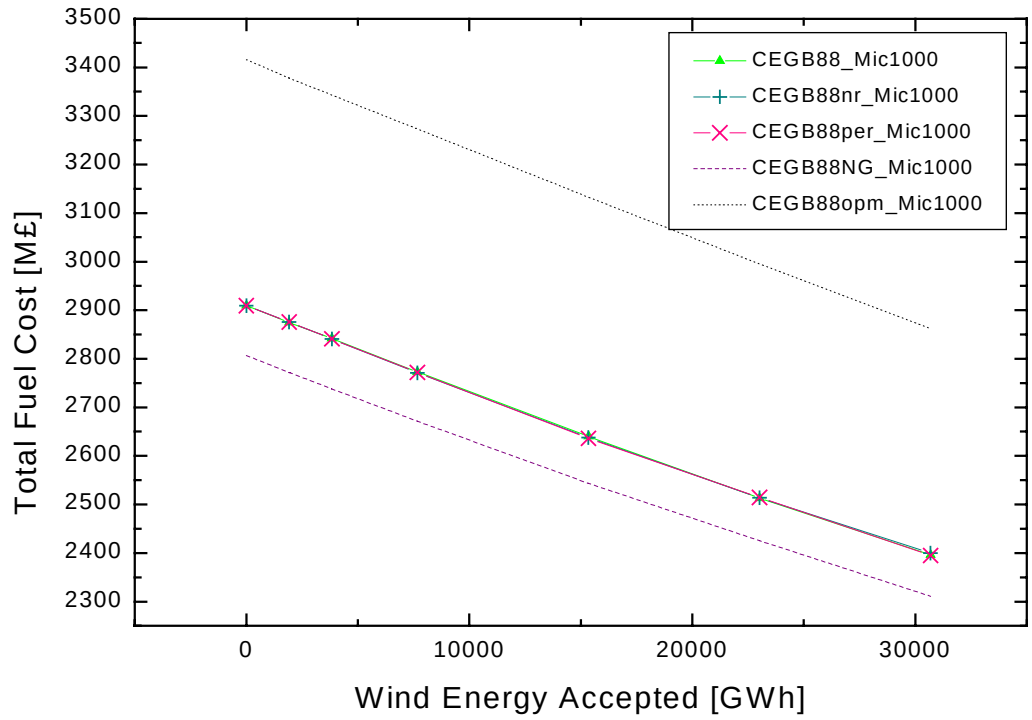

Figure 25: Different plant mixes. The lines with symbols refer to persistence, NewReference and perfect forecasting. NG, the lower graph stands for no gas turbines (all gas turbines were treated as coal plant), while opm, the upper graph, refers to the older plant mix from 1988. The turbine is in all cases a Micon 1000 kW.

Another interesting result for the value of forecasting is that in Figure 27 persistence is not the worst case. For this plot the case of Iowa was calculated [90] with the predictions from the National Weather Service (see section 2.1.3), which were used as input to the statistical WPPT model (see section 2.3.4). In this case, two effects play a major role: the NWP forecasts are done with a coarse spatial resolution and therefore are not very precise, even after using MOS [91]. In addition, the temporal resolution is only 6-hourly, in comparison to HIRLAMs usual 3 hours. Therefore, the precision was not as good as for persistence, and the economic results reflect that. However, for a penetration of up to 12% the difference is fairly small.

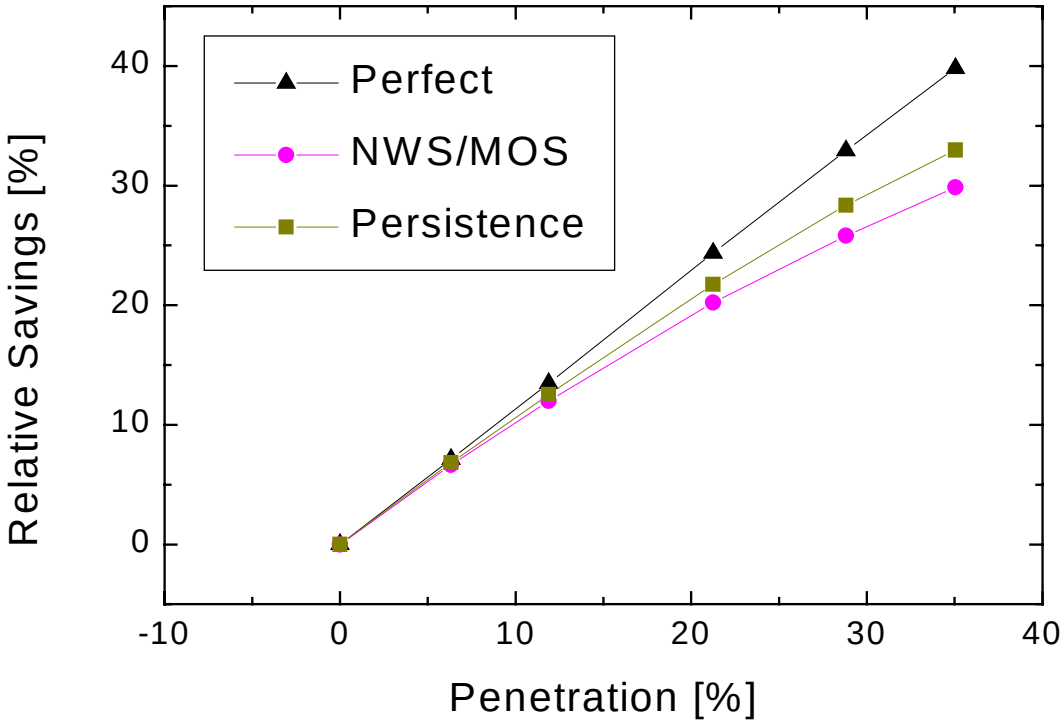

Figure 27: Relative fossil fuel savings for the case of Iowa for three forecasting options. NWS/MOS is the model from the National Weather Service. Penetration in this case is a real penetration figure. Note that the NWP approach yields worse results than persistence.

The conclusion of this chapter is that the NGM is sensitive to the type of turbine only insofar as the overall production from this turbine is different, and that the savings scale with the amount of wind energy accepted into the grid. Another important result is that for various grids, the difference between the simple persistence forecasting method and perfect forecasting only appears at rather high penetrations. Actually, the difference in forecasting precision allows more wind energy to be utilised for better forecasts, leading to higher savings of money and emissions. This effect only appears at high penetrations since it takes significant amounts of wind energy to exceed the difference between the minimum load and the base load generation.

Therefore, the analysis of the difference in forecasting options will not play a major role in the next chapters.

6. Smoothing of Distributed Wind Power Generation

6.1 Cross-correlations

The integrability of wind energy into an electrical grid depends heavily on the variations in the wind power generation. While one site (and also a small region) quite frequently logs zero and full generation, a better distributed generation will have a smoother generation profile. To which extent this is the case for Europe shall be shown here. Typical weather patterns in Europe are only about 1500 km in extent, therefore we would assume that spreading out wind energy generation further than that would be beneficial from a smoothing point of view. This chapter will deal with the smoothing effects derived from the analysis of the time series themselves, while the next chapter will assess the benefits of smoothing by simulation of the European electrical grid.

The simplest test is to investigate the cross-correlation (see Appendix) between two stations. A small value for the cross-correlation coefficient means that the time series add up to a smoother time series, while time series with high cross-correlation coefficients just add their variability.

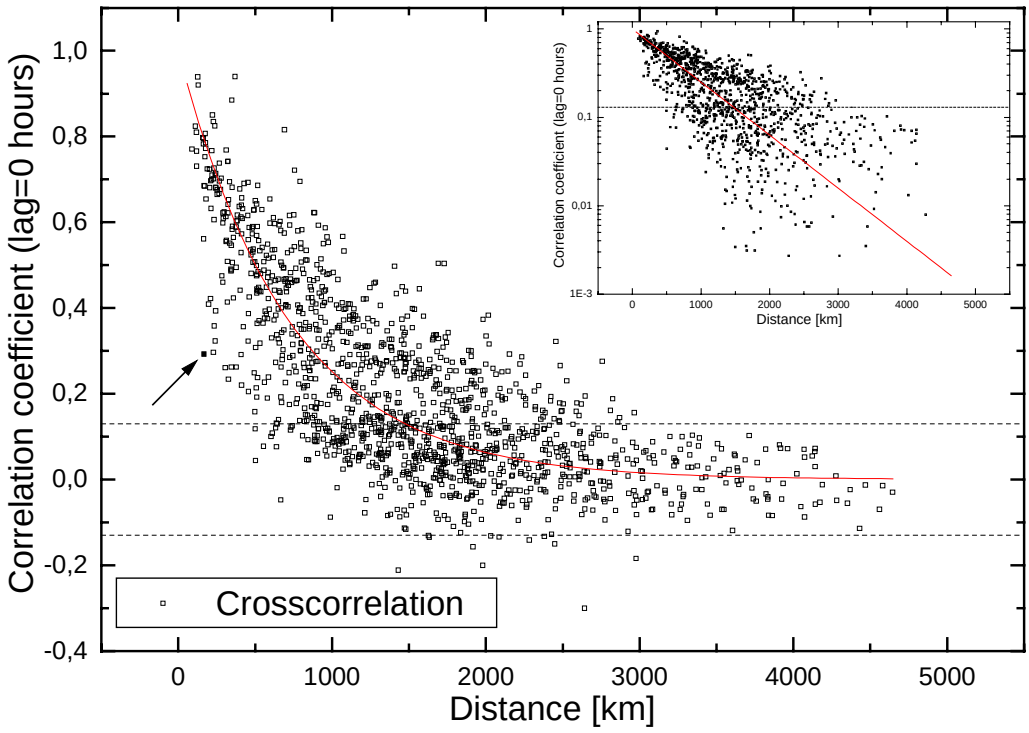

Figure 28: Correlation coefficient for every pair of stations at lag=0 hours. The dashed lines at ±0.13 are only used to guide the eye. The arrow points to the pair of stations on Sardinia/IT (see text). In the inset, the same plot is scaled logarithmically. The solid line is an exponential fit *exp(-Distance/D)*, with *D* being 723km.

Figure 28 shows the correlations for all pairs of farms with their respective distances. Care was taken to only include results, where enough concurrent data points were available (20% or more of the overall data points). While short distances give the highest correlations, a short distance does not necessarily mean that the time series are correlated. Local effects can actually lead to a significant decoupling of the time series [92]. The best example for this is the point represented by the pair Alghero/IT-Cagliari/IT,

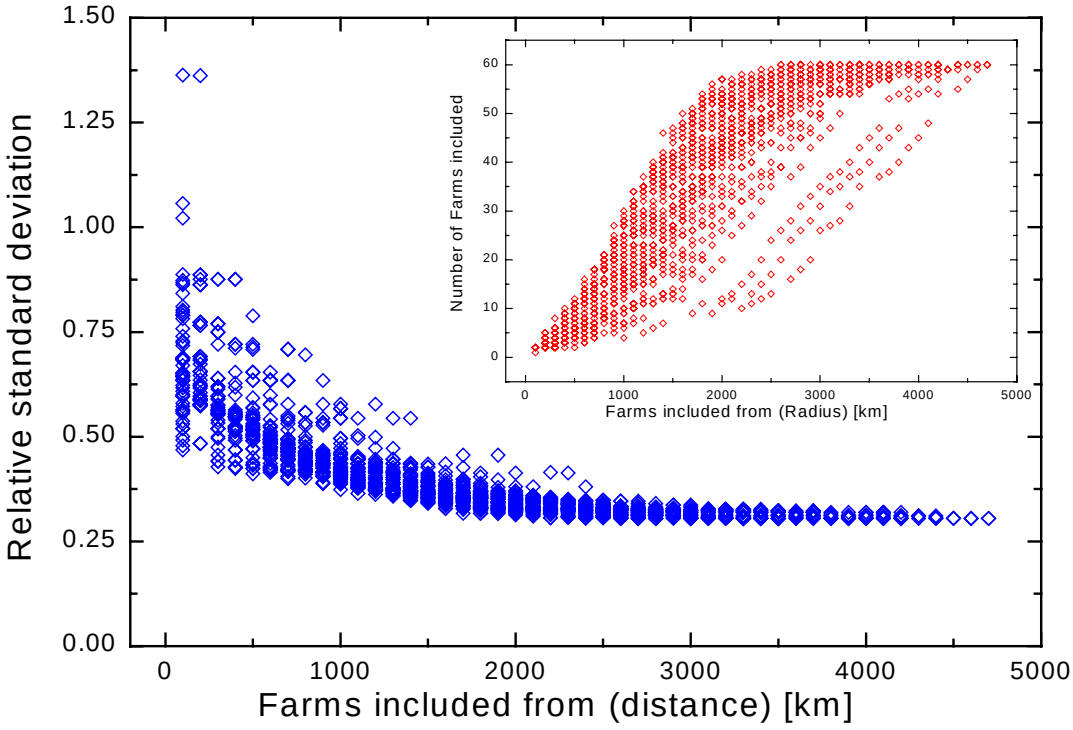

Figure 29: Relative standard deviation of the time series resulting from combining all available stations within a circle of radius R around any one station. Relative refers to the standard deviation divided by the mean of the time series. In the inset: the number of farms included for a given radius.

where the cross-correlation already for a distance of 170 km drops to 0.29. These stations are in the northwest and south of the island of Sardinia in the Mediterranean, and hence have rather different microclimates. The low wind speeds at both stations (2.9 and 3.9 m/s mean) point to local influences dominating the wind.

For longer distances, the result is as expected: the correlation is very small. Interestingly, in some cases the time series are even somewhat anticorrelated, meaning that a high wind speed at one station often coincides with a low wind speed at the other. The two pairs with the most negative correlation Roches Point/IE-Lisboa/PT with -0.21 at 1430km distance and Zaragoza/ES-Naxos/GR with -0.18 at 2977km. It is also easy to see that the average correlation decreases with distance. Hence spreading out the wind power generators should give a less variable resource. The band between 0.13 and -0.13 seems to indicate no more systematic correlation of the correlation coefficient with distance, since the numbers are equally distributed around zero. However, a slight decrease in the absolute value of the correlation is still visible. The cross-correlation data enters the band at 450 km, and all data is contained in the band at 3100 km. These numbers lead to a reduced variability of the resource, since the standard deviation of the sum of N time series is given as:

$$\sigma_{sum}^{\ 2} = \frac{1}{N^2} \sum_i \sum_j \sigma_i \sigma_j corr_{ij} \qquad (15)$$

with σ_i and σ_j being the respective standard deviations of the time series.

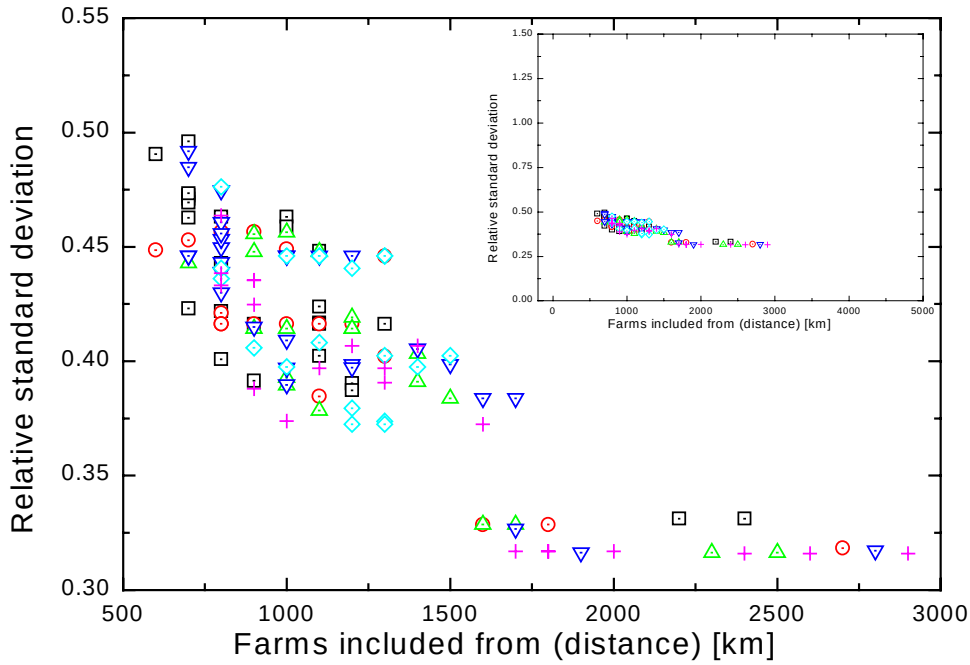

Figure 30: Standard Deviation as in the main graph of Figure 29, but this time only for data sets containing between 15 and 20 stations. Different symbols refer to different numbers of stations included for the averaging. The graph in the inset is identical to the larger graph, but scaled as in Figure 29. No reasonable correlation between symbols and standard deviation can be found.

The exponential fit *exp(-Distance/D)* shown has a decay parameter *D* of 723 km. Calculating the cross-correlations for the resulting power time series, *D* drops to 641 km. Figure 28 shows a weaker decay than other published data on cross-correlations[a]: Kaltschmitt [118] shows data with a cross-correlation coefficient of around 0.5 for distances as low as 20 km, all the way through to 180 km distance. Note that this is in fairly complex terrain, in the German *Mittelgebirge*. Steinberger-Willms [127] finds a decay parameter of 375 km for 5 sites of the DWD. However, in the calculation of the exponential fit one pair was left out, which seemed to significantly enlarge this distance. Both are hourly mean values, which should be more correlated than the 10-min means presented here. Landberg *et al* [93] show a decrease with distance for Danish stations with a decay parameter of 500 km. Their data were 10-min averages. On the influence of averaging time on the correlation decay, see Ernst [94]. A reason for the relatively weak decay could be that many stations are at the Atlantic/North Sea coast, and therefore are subject to similar weather patterns.

6.2 Smoothing of spatially averaged time series

This previous analysis was for pairs of farms only. To assess the smoothing effects from the spatial averaging of many time series, they were combined from all farms within a certain radius and the

[a] For the horizontal coherence of the wind field on scales <3km, see *eg* Schlez, W. and D. Infield: *Horizontal, Two Point Coherence for Separations Greater than the Measurement Height*. Boundary Layer Meterology **87** (1998), pp. 459-480, and references cited therein

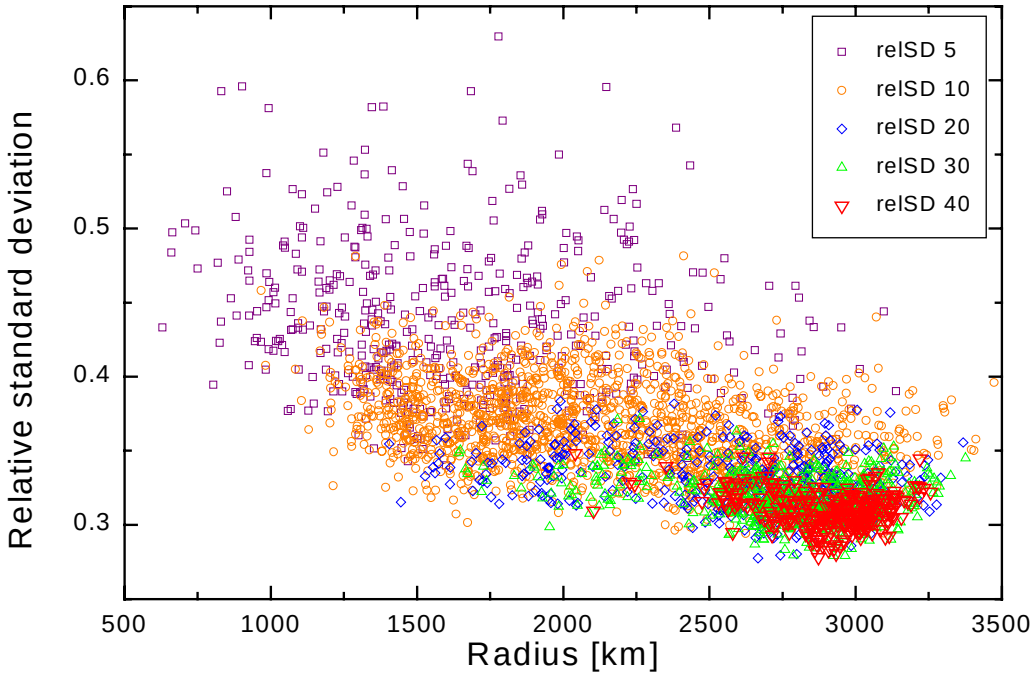

Figure 31: Relative standard deviation of time series containing 5, 10, 20, 30 and 40 randomly selected sites.

standard deviation of this resultant time series was calculated. This can be found in Figure 29. The algorithm was as follows:

At every station, an averaged wind speed time series was calculated, which included the time series of every other station within a circle around the station with radius R. The radius R was then varied in steps of 100 km. Care was taken to only include unique combinations of stations for the final plot. For every unique combination, if there was the possibility to reach the same combination from various stations, the smallest radius R was chosen as the radius for inclusion in the plot. Note that at the outside borders of the domain, fewer farms are included in the same circle, since the circles were centred on each station. This also explains the behaviour of the inset: four single lines are clearly visible, for which the number of stations included is much lower at the same radius than for most of the farms. These four stations are on the very border of the domain, in Greece.

This also shows that the time series resulting from combining many farms in a large area is considerably smoother than a single time series. The tail is dominated by the discretisation of the input time series. Another explanation for the decay with distance could be that the higher the radius chosen, the more time series were averaged. Naturally, for a larger radius more of the met stations are within the circle, hence the averaging is done including more stations, as can be seen in the inset in Figure 29.

To allow for this effect, in Figure 30 only averaged time series from a combination of between 15 and 20 stations were taken into account. Here, no real trends other than the distance dependency are noticeable; hence, the reduction of standard deviation in Figure 29 must be an effect of the distance.

A slightly different picture emerges when we change the algorithm to analyse the time series resulting from a random assortment of N number of farms. In Figure 31, a given number of farms were randomly selected and the averaged time series was calculated. The radius here refers to the maximum distance of any one farm from the centre of gravity of the farms. This was given by the mean latitude and longitude

of the farms. Hence the distribution of farms in the given circle was more uniform than in the circle defined by the radius from Figure 29, where in the case of the Greek farms only a narrow angle of the full circle actually contained farms. To find a reasonable measure for the radius is difficult. As one could see from Figure 28, the maximum distance between any two farms was 4646 km, so the radius should not exceed half of this distance by a lot. The reason why beyond this distance the relative standard deviation of the combined time series increases again is due to the algorithm used for the radius. If there are many farms in one area, and just one far away, then the average correlation between the farms is relatively high, and the centre of gravity is close to the cluster of farms, even though the maximum distance increases. Taking the average distance between the farms does not help, either.
Nonetheless, two effects are clear in Figure 31: the variation of relative standard deviations gets lower with the number of farms included, and there is a trend to lower standard deviation with increasing distance between the farms.

6.3 Properties of the averaged time series

In the next chapter, the effects of the smoothing of time series on the electrical grid of Europe will be examined. To this end, three different time series are compared (see chapter 3.6). To assess the merits of a very good farm, but without the smoothing effects, the series from Malin Head/IE was used. The other two are similar in their set-up: the series called Average is an average of all time series that were available, even those lacking a significant amount of data. Since the production at some sites was deemed not economically viable, a third series was done, largely drawing from the same area, but taking only series with more than 2000 full load hours into account. The rationale behind this was that these would be the sites that first see investment in wind energy. This time series will be called Selection in the following.
Below is a table with the main parameters for the three wind power data sets used:

Table 3: Properties of the European time series. σ is the Standard Deviation of the resulting time series, σ_{rel} the standard deviation divided by the mean, FLH are the Full Load Hours, and LF is the Load Factor. Mean, Min, Max and σ are in [kW].

	Mean	Min	Max	σ	σ_{rel} [%]	FLH [h]	LF [%]
Average:	1347	94	4085	773	0.57	1934	22.1
Selection:	1850	19	5271	1055	0.57	2657	30.3
Malin Head:	2646	0	6096	2202	0.83	3800	43.4

Note here that the Malin Head time series has by far the highest absolute standard deviation, but also the highest mean. It also shows a higher relative variation. The relative standard deviation for the Average and Selection series is identical. This means that the smoothing effects for both series are very similar, while the difference for the Selection time series is the higher production during that year. In that way, the differences between the Selection and Average series can be taken as a measure for the sensitivity to variations in power generation for the same installed capacity, leaving the relative variation the same. Actually, the variation is slightly different, since most of the good sites were close to the sea, where predominantly the wind systems from the Atlantic were prevailing. In the Average series, more inland

sites are taken into account, where the local effects play a larger role. Note also that due to the higher winds in the Selection series, these sites have higher weight in the Average series.

Here are some properties of the European Average wind profile: Maximum power generated was 4086 kW on December 26 1990 at 1100 hours, minimum was 94 kW on October 22, 0100 hours[a]. It is also worth noting that neither the full rated capacity nor zero rated output occurs during the year in question.

An important aspect for the ease of integration of varying energy sources is the rate of variation. *Eg*, hydropower has only been able to achieve the good status it has among utilities, because the time scales of variations are in the order of days or weeks, and hence easy to adapt to and much longer than most of the other time scales a utility has to deal with. Load variations are on a time scale of minutes or hours[b], power plant failures have time scales of seconds or minutes, and getting new plant on-line takes hours. Only building new plant has much longer time scales, with anything between some months for a smaller gas turbine to many years for a nuclear station. Hydro power, having mostly only seasonal variation, is well suited to the task of providing a customers load. Wind power, on the other hand, will need some serious averaging effects to make it to the status of firm power. Again, the distribution of generation over a large area comes to the rescue.

Figure 32 and Figure 33 illustrate this fact. The graphs are to be read as follows: on the x-axis, the change from one hour to the next hour is plotted (in Figure 33 it is the change on the time scale the NGM is looking at, namely 8 hours), while on the y-axis the probability density for this change can be read. This *eg* means that there is a 37% probability that the next hours output of the European average is within 1.6% of its value now, while there is a 1.6% (cumulative) probability that the output from Malin Head will decrease by 20% or more. These percentages are calculated against the installed capacity, so small changes at low levels do not lead to large percentage swings. An interesting feature can be seen in the Malin Head time series, especially for the 8-hour horizon: for rather small changes in power output the probability suddenly jumps up. This can be explained with the features of the power curve. Since the wind at Malin Head varies quite frequently in the plateau part of the power curve, the changes in wind do not translate to changes in power output. In general, the variation from one hour to the next is much lower than on an 8-hour time scale. Beware that these probability densities were calculated based on one year only. Refer to Figure 11 for the variability of the autocorrelation for various years.

In this chapter, we have seen how the distribution of wind energy generation all over Europe is beneficial for the smoothing of the resulting power output. The cross-correlation decreases with distance, to values close to zero for about 3000 km distance. This leads to a reduced relative variability in the averaged time series. The time scales of the variations are significantly reduced as well. Furthermore, the minimum generation averaged over all of Europe is about 1.5% of the installed capacity.

---

[a] In fact, maximum generation was 4415 kW at 1200 hours on December 19, 1991, but since the NGM only takes one year as an input, the last December was omitted.

[b] A quite prominent example was during the Euro'96 football championships, when in England after the semifinal England-Germany a sudden increase of about 1 GW demand could be logged within some 20 min. Germany had won in the penalties, and everyone needed a tea to calm down, most of which were brewed with an electrical kettle.

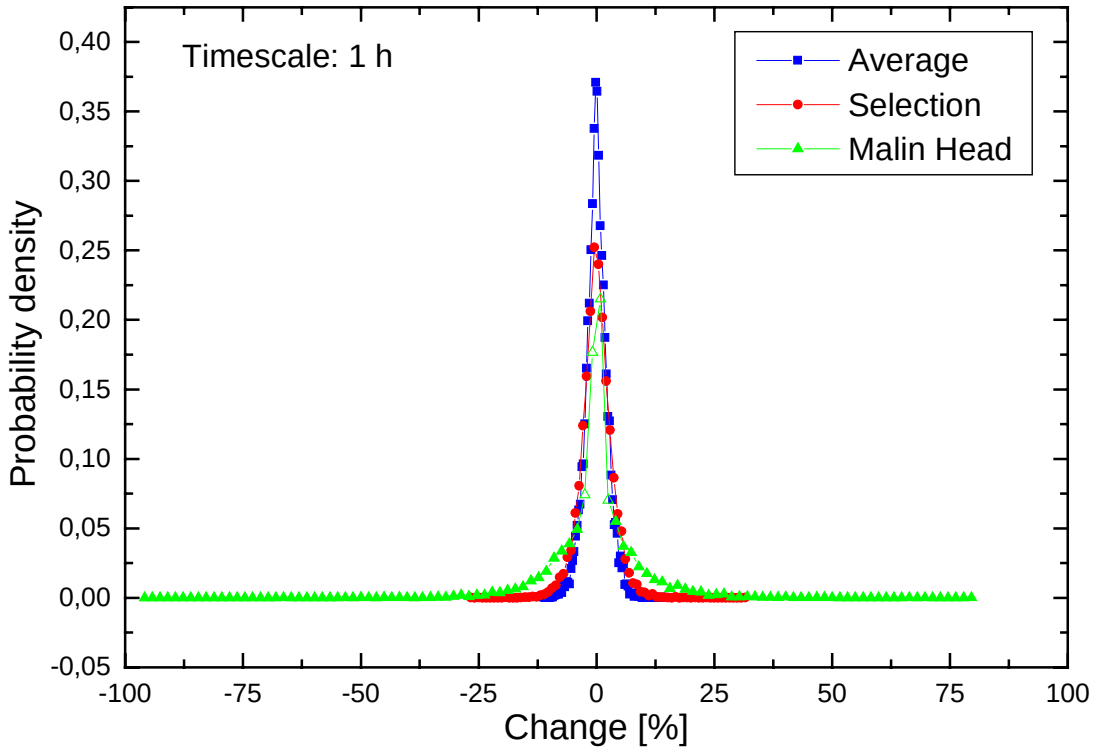

Figure 32: Time scale of the variations of the three European time series. The graph illustrates changes from one hour to the next. The x-axis is given as a percentage of the installed capacity.

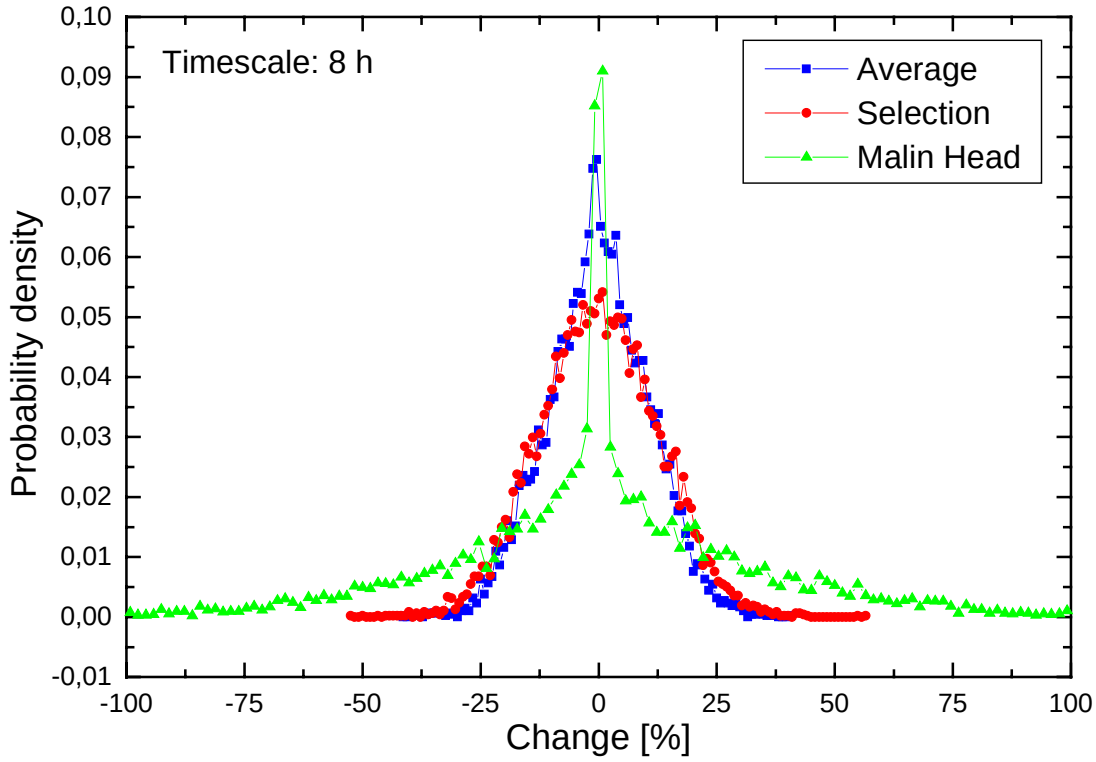

Figure 33: Time scale of the variations of the three European time series. The values have an eight-hour interval.

7. Capacity effects of wind energy in Europe

7.1 Introduction

In the chapters leading to this page, we have seen that the resource can be predicted quite well, even though predictions only make a difference for high penetrations, and that the distribution of generation smoothes the output. A nice side effect of the smoothing is that a less varying time series is also easier to predict. In this chapter we will make use of all things learned so far. Let me start by briefly rehashing the motivation.

Wind energy is currently the fastest growing energy source in a number of countries [95]. Nevertheless, even though it reached quite sizeable local penetrations, the overall impact of wind energy in Europe has not been really visible. This chapter will try to answer some of the questions arising from higher penetration of wind energy in Europe.

Two questions are rather strongly disputed in the utility world: How can we integrate a lot of this variable resource, and how much conventional power plant is no longer needed due to the installation of wind power? These questions are important for the standing of wind energy and, together with the price of wind energy, will determine whether wind energy is seen as a viable option for large-scale electricity generation. Up to now, most people (and especially most utilities) believe that wind energy can not contribute to the firm generating capacity of the grid[a]. While this can be so for a single turbine which can be running at full output or not at all, this is no longer true when looking at many turbines distributed over a large area, *eg* Europe. Parties interested in this answer could be utilities with the task of running a high-penetration grid, or politicians trying to decide on a fair payment for wind power.

The ability of a grid to accommodate a variable power source is highly dependent on the relative time scales involved. If the grid operator has many fast power plants available, then more variable generation can be dealt with. Fast in this context refers to the time scales involved: while one wind turbine can have gradients from 100% to 0% within a few seconds (in the case of a safety stop due to a passing storm), this does not hold true for dispersed turbines, where the storm front will hit at different times and in different magnitude. Hence, the time scales of wind energy get longer with a higher dispersion of the generation. Longer time scales mean better ease of integration into the existing power plant mix. Connected to this is the question of whether wind power generation can be counted upon at all times, or whether conventional power plants have to be on standby, increasing the cost for the grid operator.

Some other answers to less disputed questions will be given along the way, such as the potential for wind energy in the existing grid, the reasonable penetration for wind energy, and the economic savings attainable through the use of wind energy.

[a] This refers to letters to the editor in the Frankfurter Allgemeine Zeitung over the last few years. Every other time the topic of wind energy is in the news, someone will write in and claim that it is an unreliable resource, and to account for the unreliability there has to be the same amount of spinning reserve in the grid. This spinning reserve would then have to be taken as the ecological and economical backpack of wind energy.

7.2 Definitions and terminology

In this section, some of the terms used in the rest of the chapter will be introduced.

Installed Capacity refers to the nameplate capacity of the turbine multiplied by the number of turbines.

Penetration is the percentage of installed wind energy capacity in the grid. Occasionally, penetration is also used to denote the percentage of load served by wind energy. Note that in the following the additionally installed capacity is used, scaled as a percentage of the existing grid, as opposed to a proper penetration figure; the reason is an easier batch routine set-up. When not discussing values, the terms penetration and additionally installed wind capacity might be used interchangeably.

Load is here used for the demand on the electricity grid.

Load Factor (LF, also known as **Capacity Factor**) is the percentage of power production as a fraction of the nameplate capacity of the wind energy conversion system. This can be the instantaneous value, but often will be the yearly mean. The latter can also be expressed as **Full Load Hours** (FLH) via a multiplication with 8760. Typical values are 20-30%, or 1500-3000 FLH respectively. The current mean capacity factor is 0.23-0.25 (or 1750-2000 FLH) [95].

The easier one to assess of the capacity effects of wind energy in the grid is the **firm capacity**. This is the fraction of installed wind capacity that either is online at all times or with a probability similar to the availability of a fossil fuel power plant. Fossil fuel plant availabilities seem to have large scatter in the literature. The CEGB used 79-92% [104], while Bernow *et al* used 84% [101], quoting the US national average for the forced outage rate as 12.4% and the maintenance outage rate as 13.6%. The firm capacity might also be expressed as an absolute value, but similarly to the LF, the fraction is usually more meaningful. This value can be assessed comparatively easily from pure wind data - in fact, this more or less has been done by Landberg [77]. According to this paper, wind turbines will only in 2% of all cases not produce power anywhere in Europe. However, we have seen in Table 3 that using the same data it is possible to come to different conclusions.

The **Capacity Credit** (CC) assigned to a regenerative conversion plant is the fraction of installed (regenerative) capacity by which the conventional power generation capacity can be reduced without affecting the loss of load probability [96, 97].

Loss Of Load Probability (LOLP) is the probability that a **Loss Of Load Event** (LOLE) occurs. Typically, system operators aim for 1 event in 10 years (or better, of course). For the LOLP, the match between resource and demand is decisive, as well as the response times of the existing power plants. Power supply systems with a high percentage of storage (*eg* pump storage) can accommodate higher penetrations of wind energy than supply systems consisting solely of nuclear and coal fired plants.

7.3 Previous works

While many studies on wind energy in Europe and the capacity credit of wind energy have been made, none so far has looked at a common European grid. At the end of this section, a short summary of the lessons learned will be given.

Selzer [98] showed that Europe's electricity needs could be met with wind energy a few times over. However, the integration into the grid set technical limits for the possible penetration due to the fixed regime of nuclear and cogeneration plant: at high wind energy generation and low demand, some of the

wind energy has to be discarded, if not enough storage capacity is available. Based on four national studies, he estimated that the available storage capacity facilitates 10-20% of the demand covered by wind energy. As well, with the proper operational strategies for the conventional power plants, the rate of acceptance could rise from 8% of peak demand to up to 30%.

Even from the early stages of development the capacity credit of wind energy has received attention by researchers. A short overview over early results can be found in Diesendorf *et al* [99]. Let me quote a few main points also found in other publications: for small contributions of wind energy (<5% of total demand), wind energy's capacity credit is roughly equal to the average wind power. (However, this is only true on average - see chapter 0.) At large penetrations (>40% generation), the capacity credit tends towards a constant value determined by the loss-of-load probability without wind energy and the probability of zero wind power. A grid composed of few large power plants attributes a higher capacity credit to wind than one composed of many small units. Dispersion of wind power can raise the capacity credit about 20% over its value for just one site. A good correlation between wind power and demand can lift the capacity credit by about 20%. Even though these results were derived with relatively simple tools, and though the actual numbers will most likely depend on a number of factors pertaining to the local circumstances, most of the results are still valid.

In a study for the Cape Cod service area, Johanson and Goldenblatt [100] concluded that it is necessary to utilise hourly wind and load data to establish the value of wind power to the utility. The monetary value to the utility of each added turbine is less than the last turbine, since every turbine replaces power plant further down in the merit order, *ie* the replaced plant has lower fuel cost than the one replaced before that. This is a quite interesting point for the economic assessment of wind power: for a given price of turbine, there is an optimum penetration of wind power in the existing grid. Reoptimizing the plant mix after the inclusion of wind power, more peaking and less base plant was found to be optimal in comparison to the case without wind power. When comparing this mode to a pure fuel saver mode, where the conventional plant still runs as spinning reserve, a larger capacity credit can be achieved, and more money is saved by the utility on fuel and investment. However, for rather small penetration the difference in value for the utility of the two modes was negligible, and only after installing a wind power capacity of 5% of the peak load, did the reoptimisation step yield higher savings.

Bernow *et al* [101] used the ELFIN model for the case of a small utility in the mid-west US, and again found the capacity credit to decrease with penetration. The percentage figures for penetration were scaled by the peak load. They explicitly analysed the benefits of adding another site for diversity of the resource. While for just one good site the capacity credit decreased from up to 100% for nearly zero penetration to 40% at 20% penetration, adding another site decreased the initial capacity credit to between 60 and 80%, but kept it above the single site CC at 20% with 45-60%.

One of the largest studies of the potential benefits in all national electricity grids in Europe was the Wind Power Penetration Study sponsored by the EU Commission. During this study, the capacity credit was assessed for each of the then 12 countries of the EU.

In the Irish study [102], the capacity credit was assessed with a LOLE method for one single farm and for a collection of farms separately. The capacity credit for the single farm saturated at about 200 MW of wind generators installed (in a grid consisting of 3800 MW generation capability), while for the

collection of sites the reduced variability allowed capacity credit increases up to about 350 MW. Again, the relative capacity credit dropped from over 30% for both options (slightly higher with the one farm for very low penetration) to 9% (collection) and 5% (single site) at 4000 MW installed wind power. An early study on wind energy in Ireland [103] concluded that the high contribution of base load plant in the grid of the ESB (43%) lead to a high proportion of wasted wind energy, since the fixed regime of the base load plant had already all the demand covered. A wind energy installation generating 15% of the total demand would waste 50%. Reducing the amount of base load plant, the losses were greatly reduced (ca 10% at the same level of installation when using only 20% base load plant). Another conclusion was that smaller generator sets for the same rotor size could lead to a greater rejection of available power as penetration increases, since the load factor of the turbines increases. For a percentage of 5% of the total demand covered by wind energy, they found a capacity factor of 52% of the installed wind capacity for turbines having rated/mean wind speed ratios of 1.5, while for the more typical case of 2 they still had a capacity credit of 24%. This dropped to 21 and 7 %, respectively, for 35% of demand covered. From a load duration curve *Ansatz* they found that the demand for peaking plant (gas turbines) would rise by 92%, while the need for base plant would decline by 39% (at 35% demand covered by wind).

For the CEGB system (England and Wales) [104], penetrations of 2, 5, 10 and 15% of wind energy netted capacity credits of 31, 25.6, 19.4 and 15.6% of the installed capacity. However, these numbers are not quite comparable with each other, since for a penetration up to 5% all installed capacity was thought to be erected on land with a capacity factor of 34%, while from that point on wind energy capacity was being built off-shore, with an accordingly higher load factor of 44.8%. Hence, for very small penetrations the capacity credit reached the average load factor. These load factors are for the high-load period in winter; the annual load factors were 25.1 and 35.1%, respectively. At that stage, Lipman *et al* [105] had already shown for the England and Wales grid that quite large amounts of wind energy (up to 30% of demand covered) could be integrated there without incurring large fuel penalties from having to cover for outages with spinning reserve.

In the Portuguese case [106], for additionally installed wind power capacities of 7.8, 18.7 and 30.6%, respectively, capacity credits of 36.5, 28.9 and 22.9% of the installed wind power capacity were found, measured against a coal fired power plant with a forced outage probability of 17.3%. These capacity credits are relatively high, which reflects the high amount of hydropower with reservoirs (2000 MW out of 6300 MW total installed capacity) in the Portuguese system. The integration of wind energy also leads to significant amounts of avoided pollutants.

Quite similar results were found in Denmark [107]: at penetrations of 5, 10 and 15%, the capacity credits were 30, 25 and 20% of the installed wind power capacity, respectively. This was compared to a previous study [108], which had found a capacity credit of 23, 16 and 11%, respectively. They attributed the difference to the "much better capacity situation" assumed in the older study.

The high end of the scale is also found in Greece [109]: for 2.5, 5, 10 and 15% penetration the capacity credit calculated was 38, 27, 20 and 17%, respectively. The reason for these high numbers is the very high assumed wind energy generation from the turbines: at these penetrations, the load factors were 49.5, 45, 41.3, and 32.3%, respectively. This decrease reflects the economic decisions of wind farm

developers to start with the most promising sites and then subsequently to spread to sites with lesser resource. An interesting by-result was that using wind energy in energy saving mode, leaving the replaced plant operational, facilitated higher monetary savings than in capacity credit mode, even including O&M of the plant and investment cost. This is because in the optimised power plant mix including wind energy, the energy savings come mainly from relatively cheap imported coal, while in the existing plant mix wind energy replaced generation from the more expensive fuel oil.

For the Netherlands [110], a previous study had shown the possibility to integrate wind power from 1000 to 1600 MW installed capacity (in a grid of about 16 GW) without recourse to storage and without significant amounts of wind power discarded. In their case, most of the energy from wind replaced energy produced in base load plant. At 1000 MW installed wind power the remaining variations in power output of minutely values even on a scale of 30 minutes are significant: a change of 500 MW has a frequency of just under 1‰. However, van Wijk *et al* [111] did a rather extensive study of the potential power production in the Netherlands, identifying potential sites and using representative meteorological data from 11 sites. They found annual load factors between 19% and 26% in a 10-year period. The maximum output attained was 94.5% of the installed capacity (barring maintenance and forced outages). Some of the losses were attributed to wake losses. They also found that the difference in wind power production between two successive hours never exceeded 40% in the 10-year period analysed. No reason for this discrepancy could be found from the two studies. Also for the Netherlands, Halberg [112] assumed a capacity credit of 20% for 5% of demand covered, while he saw it dropping to 13% at 15% contribution.

Using a PreussenElektra model, Consulectra found for Germany [113] a capacity credit of 15% at 10% penetration. However, just using the coldest days of the year, the capacity credit dropped to only 6.7%. This feels peculiar, since usually the winds in winter are strongest. They also show a generation duration curve for one site in Ostfriesland (Esens), where the curve for the coldest days shows zero generation in more than 50% of the time, while in only 30% of all times during the whole year no generation was found. Earlier, Jarass [114] had shown that a cluster of wind parks on the German coast has a significantly smaller capacity credit when feeding into the (smaller) coastal grid compared to feeding into the grid of all (Western) Germany. The same number of wind power plants distributed over all of Germany displaced even more conventional capacity: 23% at 15% of the total demand covered from wind energy.

Another result not quite fitting the overall picture comes from Spain [115]: the capacity credit is calculated as less expansion between 1993 and 2000, and is 10, 16.8, 15 and 15.6% of installed wind capacity for 1.5, 5, 10 and 15% penetration. Surprisingly, the capacity credit is not monotonously decreasing, as was the case with most other countries. The reason could be the discretisation of power plants being replaced: at 1.5% penetration, wind power replaces 1 plant of 92 MW, while at 5% it replaces 2 plant with 550 MW. Therefore, the capacity credit is dependent in their calculations on the expansion plan assumed. All the capacity replaced was newly to be built power plant running on imported coal.

The case of Italy [116] is special, since the resource is good enough for exploitation only in a few selected regions. Hence, the target penetrations were only achieved in relation to the electrical grid in

the region. As well, their value for the capacity credit was well below the ones found by the other teams (22.6% at 0.5% national penetration, corresponding to 2.5% penetration in the windy regions).

Fischedick and Kaltschmitt [117, 118] assessed the potential for wind energy in the German Land of Baden-Württemberg by using hourly data from 7 stations and spatially interpolating the wind power for every region. Taking the technical potential into account, about 5-8% of the electricity demand of the Land could be delivered by wind energy. Analysing the case of a smaller utility with a high percentage of pump storage, they found that the optimal use of wind power is made with a plant mix of small power plants embedded in a grid with large storage capability. However, their result differs from most other results in that wind power mainly replaces power plants in the medium and peak load segment, and nearly none in the base load segment. This could be connected to their rather low load factor for wind energy generation (21.2 and 15.2% for low and medium/high penetration, respectively).

High wind energy penetration was first examined for the case of island grids, where the installed capacity is typically relatively low. A sizeable proportion of wind power is therefore already achieved with few turbines. Papadopoulos *et al* [119] investigated the cases of Crete and Hios, two Greek islands. They found very high wind power gradients from one minute to the next (up to -62.2% for one farm, and -38.3% for three farms on three islands of the Cyclades group, approx. 40km of each other). They concluded that high wind penetrations in large island could be achieved, as long as the power system as a whole is designed to deal with large wind power fluctuations. A common dispatch centre and wind energy forecasts were deemed essential.

Saramourtsis *et al* [120] showed for the Greek island of Syros that the percentage of wind energy generation accepted into the grid highly depends on the maximum fraction of wind allowed into the grid at any given time. With a 10% limit, the permitted wind energy already trails off for a penetration of 3%, while with a limit of 50%, discarding wind energy only starts significantly at more than 10% penetration. This low attainable penetration might have two reasons: for one, the wind power generation at the different turbine sites was assumed to be identical for all wind turbines, and the available generation consisted of just 5 large Diesel sets, which might be oversized (for reasons of reliability) in relation to the actual load.

Hansen and Tande [121] showed that the modelling approach used for low penetration studies is valid as well at higher penetration. With this aim they did a sensitivity analysis for 2.5 and 25% penetration for the case of Praia, the capital of the Cap Verde islands, and showed that most parameters did not heavily influence the levelised production cost for the whole energy system, even though the variation width was slightly higher with higher penetration. The only factor to strongly influence the economics of the wind power development was the average wind speed at the site in question. The also claim [122] that for small numbers of wind turbines, the assigned capacity credit has to be reduced due to the fluctuations in output.

Using Markov chain modelling, Torre *et al* [123] set by the limits of wind energy penetration for Corsica/FR to 30%. However, this was highly dependent on the ratio between installed capacity and peak load. The 30% limit was reached at 70% ratio, while for 80 and 90% nearly no wind energy could be integrated, and for lower ratios, wind energy could easily be integrated up to much higher penetrations.

The network strengthening properties of distributed generation are especially apparent in remote areas with weak networks and high resources. In rural India *eg* [124], building high power grids is rather expensive, while area for renewable generation is not an issue.

Lakkoju [125] generated joint distributed probability density functions for the combined generation of offshore wind power with onshore wave power devices. The combination yielded a higher probability for generation at medium and high output than the two generation options alone.

The integration of variable sources of energy can also be tackled from the demand side as demand side management. When the demand can react according to the wind power offered, more of the variable source can be used to supply demand. Simple appliances on the residential customers side where the demand can be shifted for at least a few hours include washing machines, dish washers, freezers and hot water preparation. Most other appliances have only very limited shifting potential (read: shorter than one hour), such as refrigerators or space heating.

The influence of variable pricing on residential customers was assessed in a Finnish experiment [126]. The idea is that if customers are getting clear signals on the actual price of electricity, they can shift some of their loads from peak hours to off-peak hours. They found that only about a quarter of all customers reacted strongly on the variable pricing, hence the potential seems to be limited.

Having access to electricity at varying market prices, the savings possible for electrical cars were estimated by Nielsen *et al* [4] to be in the range of 11 €/a. Therefore, the incentive given is very low.

One study deserves to be singled out, since most of the problems laid out in this work were already tackled therein: the dissertation of Robert Steinberger-Willms [127]. In comparison to my study, he also looked at an energy supply from solar energy, including solar power from solar thermal power plants in northern Africa transported via HVDC. For the wind part he worked with just 5 hourly time series from northern Germany. He could show that even with just the 5 time series together, the frequency spectrum was nearly one order of magnitude lower than for just one park at frequencies of 1/6h. This means that the spatial averaging especially takes out the short-time variations in the resulting time series. He could show that due to the stochastic nature of the wind generation (on a time scale comparable with the load variations, that is, intraday), high penetrations of wind power replace base load plant, while the good match between peak load and peak generation has solar power replacing predominantly peak load plant. The available energy storage does not need to be large to be of great benefit for a renewable generation system: even 3 hours worth of storage[a] allow renewable energy to replace much more conventionally generated electricity than without, and the additional benefits already trail off significantly at 12 hours. However, this is calculated for renewable penetrations of up to 1, using 10 times as much energy from solar than from wind power[b]. The surplus energy from renewable generators was highly dependent on the amount of base load plant. At high penetrations, much energy had to be discarded when 25% of the power plant mix was base load running on a fixed regime; this lead to the conclusion that high penetrations of renewables demanded a reoptimisation of the existing conventional power plant mix.

[a] Read: a storage system being able to cover the average demand for three hours.

[b] This was modelled after the resource available in Germany. A penetration of 1 means a system completely running on wind and solar energy.

Let me summarise this: Wind energy has a capacity credit. This is depending (among other things) on the load factor and the penetration. The capacity credit tends to decrease from approximately the load factor for small penetrations to some 10-15% at high penetrations. It is highly dependent on the electrical system used for comparison, especially the amount of storage possibility and the match of load and demand. A proper assessment needs hourly load and wind power time series. Wind energy predominantly replaces base load plant, since more flexibility in the system is needed to accommodate wind energy. Base load plant cannot be regulated, and is therefore detrimental to wind power integrability.

In the following, a common European grid will be investigated, some problems in the capacity credit assessment with time domain models will be illuminated, and some ways to circumvent them.

7.4 Benefits of distributed production

The time series analysed in chapter 6.3 will now be used as input for the National Grid Model. The inner workings of the model have been laid out in section 2.7. Here a short summary of the analysis performed by the model. The European load time series, plant mix and fuel cost discussed in 3.7 is used as input, together with the three time series discussed in the previous chapter. These were a European Average production, another time series of distributed generation, but only using a Selection of sites with good resource, and the single time series with the highest generation (Malin Head). The scheduling algorithm in the NGM tries to match this demand time series hour for hour with the given power plants and the available wind energy. The wind time series are multiplied with a factor yielding a multiple of 2745 MW installed capacity, which is somewhat below 1% of the installed capacity of the conventional grid. The conventional plant mix consists of some 3000 power stations with an installed base capacity of close to 350 GW.

The National Grid Model essentially calculates the minimum spinning reserve requirement (using minimum fossil fuel cost as the quality function) so that no loss of load events occur in a given time period. This minimisation of fossil fuel use also leads to other environmental benefits, such as less air pollution and less CO_2 emissions. In order to assess the benefits of wind energy for the grid, wind energy is added from 0 to about 80% additionally installed capacity. This corresponds to 45% penetration. The total cost of fossil fuel for one year is chosen as an output variable. In the European grid analysed here, the total cost for fossil fuel for one year without any wind energy is 11.63 G€.

The main benefit of wind energy in the grid is that it supplants fossil fuels. However, while at small penetrations of wind energy this benefit scales practically linearly with the generated wind power production, two additional effects appear at higher penetrations. Small penetration here is where the variations of the wind energy are the same magnitude as the load uncertainty (about 1.5% of the load).

The first of these effects (Figure 34) leads to significant amounts of discarded wind energy at high penetrations. The ability of the grid to accommodate the wind power production is limited by the power plants that can not reasonably be regulated, like nuclear plant, or have preference over wind energy (like hydropower), or are insufficiently flexible like thermal plant at minimum load factor. If in this case some fraction of wind energy already is enough to cover the demand, the remainder has to be discarded.

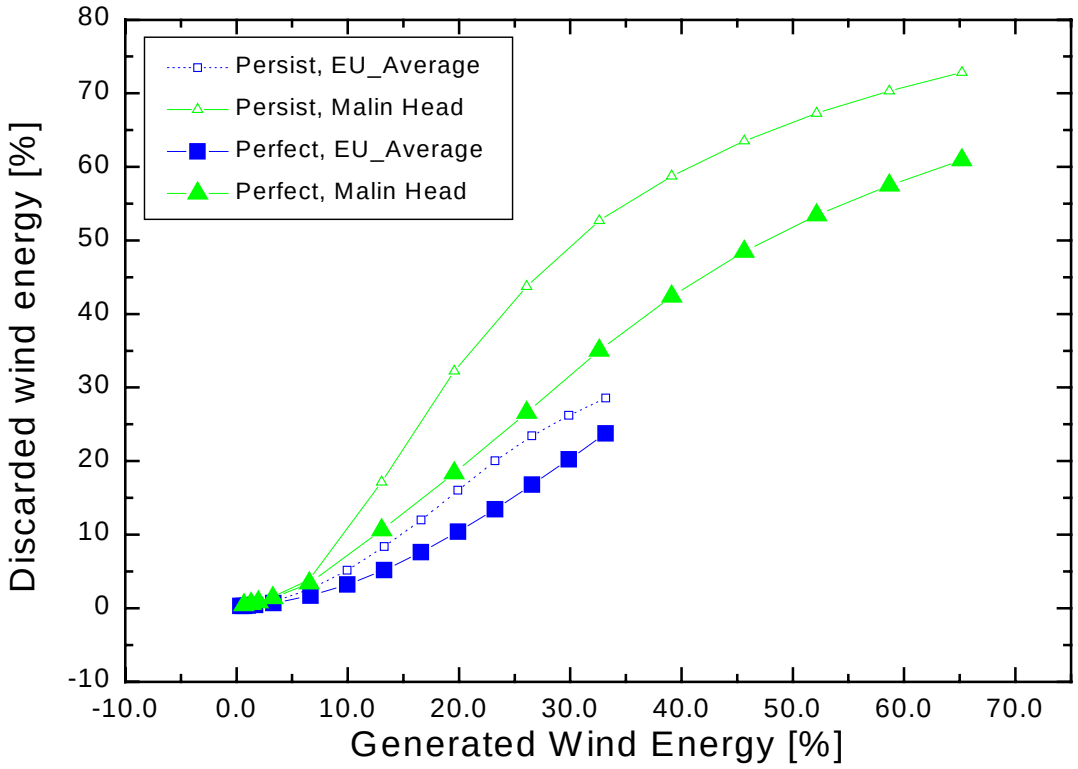

Figure 34: Discarded wind energy production. This figure shows the percentage of wind energy that is not accepted into the grid, as a function of the generated wind energy, which is scaled as a percentage of the total demand. Persistence and Perfect are the two forecasting options. The Selection series was omitted; it behaves mainly like the Average series, though it is longer.

It is easy to see that for small penetrations, nearly no wind energy has to be discarded, while at higher penetrations the percentage of discarded wind energy rises strongly. This result also shows that wind energy can be integrated to about 15% of covered demand with just 10% of wind energy wasted due to the assumptions built into the model. This is only valid for the averaged time series - Malin Head already wastes 10% at about 10% of demand covered. For perfect forecasting the result becomes even better: nearly 20% of the total electricity demand can be covered before 10% of the wind power production has to be dumped. Using perfect forecasting for Malin Head, the result is 12.5%. The difference in the forecasting algorithms is due to a slightly different effect than the reason given before for discarding wind energy. Not only the base load plant is regarded as fixed; if steam plant is already regulated down to its minimum in the scheduling process for the next hour, it cannot be dropped from service, if the wind is exceeding the forecast. Therefore, the excess wind power has to be discarded, too. Discarding 10% of the generated wind energy may sound like a lot - but it is not, due to the assumptions made in the NGM. In this case, the unmodelled reservoir-fed hydropower will take most of the energy provided. Of the about 100 GW installed in Europe, about half is hydropower generators connected to some kind of reservoir. At the very least, it is capable of holding back some generation for a few hours, though typically the time scale should be weeks. Therefore, most of the energy here would not be produced at once by hydropower, but deferred for later use. Figure 34 tells us that good forecasting combined with a low variance wind production leads to a better integrable resource, while high variability and bad forecasting leads to much wasted wind energy. We can also see from the graph that

Malin Head has both the highest generation and the highest percentage of discarded energy. Taking both effects together, this leads to Malin Head not serving most energy to the grid - it is the Selection time series, which has higher generation than the Average series at the same installed capacity, but nearly the same smoothing effects due to the distribution of generation. This can *eg* be seen in the inset of Figure 36: the Selection series goes furthest to the right, also facilitating the largest fossil fuel savings.

The other effect leading to less than linear fossil fuel savings is less easy to understand. This is shown in Figure 35. As more and more wind energy is added to the grid, the spinning reserve (see section 2.7) has to increase, since the extra insecurity of the wind energy generation has to be accounted for in order not to have a loss-of-load-event. This spinning reserve however has to be delivered by less and less fossil fuel power plants (as we will see in the next chapter). Herewith it leads to a lower average load factor for the remaining fossil fuel plants. Keep in mind that the minimum part-loading of fossil power stations assumed in the NGM is 50%.

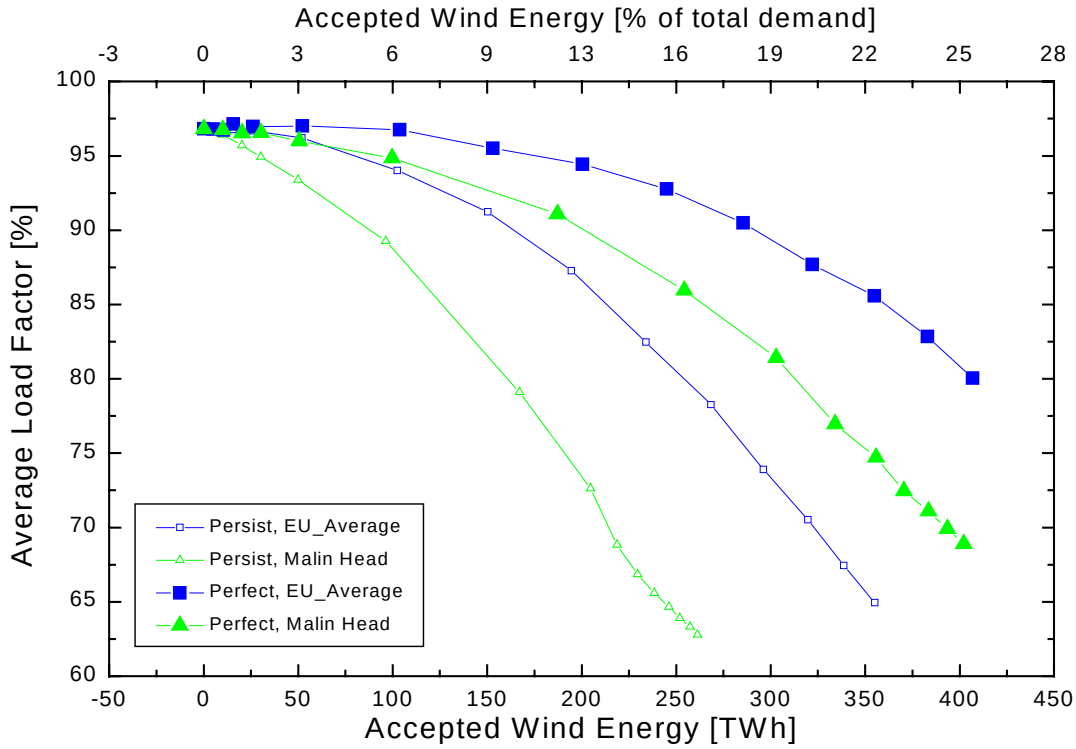

Figure 35: Average load factor of the remaining fossil fuel plant.

In Figure 36 we see that for small penetrations the possible savings correlate with the amount of wind energy produced and fed into the network. The shape of the graphs in Figure 36 is determined by the two effects pointed out in the foregoing paragraphs, the discarding of wind energy, and the operating penalty due to reduced load factor of the remaining fossil fuel plant. Actually, at high penetrations the fuel savings correlate with low variability of the time series and high forecast accuracy - the highest savings for very high penetrations are attainable with a mixture of perfect forecasting and high wind energy generation. All the data points (except the first four) of the different graphs are equidistant in installed wind capacity. The saturation effects for high variability of the input, coupled with bad forecasting (Persist, Malin Head), are clearly visible, even though in no case it reaches full saturation. Note that the spread between the forecasting methods is larger for more variable wind energy generation. In the inset in Figure 36 we see that indeed the fossil fuel savings scale with the amount of

wind energy admitted to the network. If we now look up in the inset the savings possible at 20% of demand covered for perfect forecasting and the Average series, we find that about 7 G€ fuel cost for fossil fuels can be saved by wind energy. This corresponds to 60% of the total fossil fuel cost without wind energy. Using these figures, we can assess the value of wind power at that penetration in pure fuel saver mode to 7G€/(0.2*1600TWh)≈2.2€c/kWh.

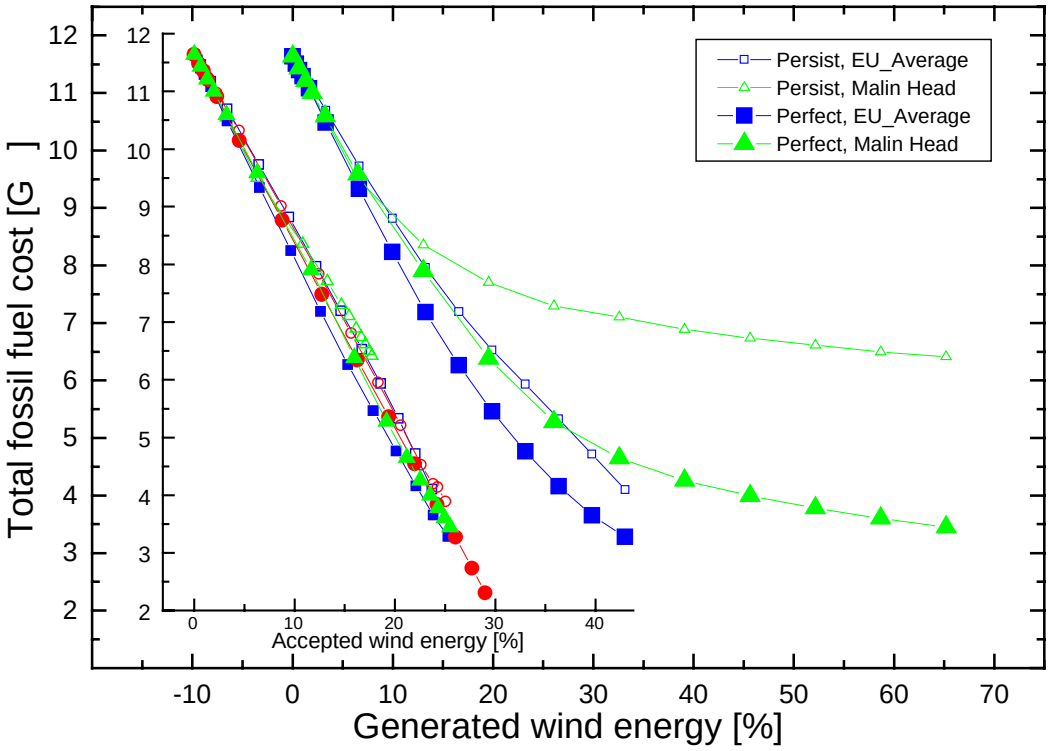

Figure 36: Fossil fuel cost for different wind energy production. The x-axis denotes wind energy produced by the simulated turbines, scaled to the total load. In the inset, the same graph is shown as a function of the amount of wind energy, which is actually admitted to the system. The red lines in the inset come from the Selection time series.

7.5 Replaced fossil fuel capacity - the single event problem

For the assessment of a capacity credit, the replacement of conventional power plant by wind power should be investigated. Since the NGM does not take forced or scheduled outages into account, a proper capacity credit assessment based on the calculation of LOLP is not possible with the NGM. It is however possible to ask the NGM how many fossil fuel plants were not used during the run. In the following, this will be called the unused fossil fuel capacity. If we assume that with the high number of conventional plant in the grid, the outages of single plants are averaging each other out, then we can use the unused fossil fuel capacity as at least some measure of the capacity credit. The unused fossil fuel capacity is the capacity used at zero penetration minus the capacity used at higher penetration. Since the unused capacity conceptually corresponds to the firm capacity of the replaced fossil fuel plant, and not to the installed capacity [see *eg* 102], a first order approximation of the capacity credit should then be the unused_fossil_fuel_capacity/reference_plant_availability. Milligan [128] and Jarass [114] point out the difference between a proper capacity credit assessment and this shortcut, since also the effective firm capacity of a fossil fuel plant is dependent on the details of the other already installed power stations in

the grid. In the literature outage probabilities are hard to find, and the values have a quite wide scatter from about 5 to 25% [104, 129]. Bernow quotes a national average in the US for forced outages of 12.4% [101]. Therefore, whatever is assumed as a forced outage probability of the plant wind energy is replacing, the unused capacity shown here should be increased accordingly. In Figure 37 the unused steam and gas turbine capacity is given.

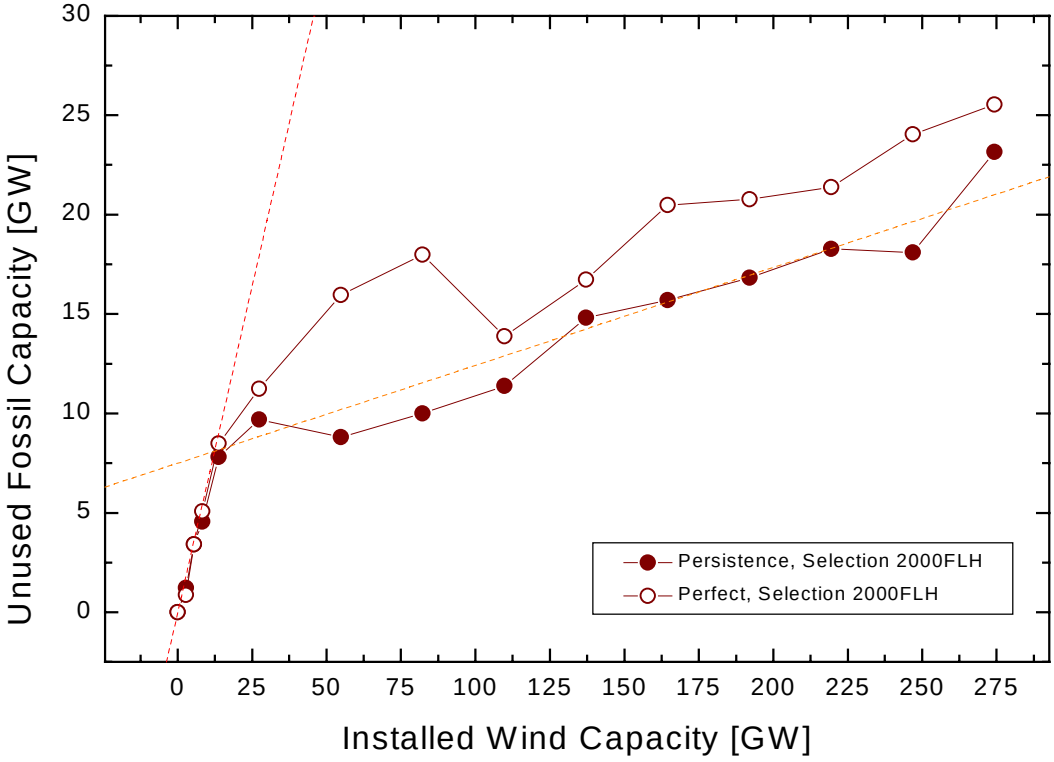

Figure 37: Unused fossil fuel capacity for the Selection time series. The dashed lines are just to guide the eye: the first having a slope of 60%, while the other is 15%.

For reasons of clarity, only the result for the probable wind sites is given here (the Selection time series), but the result for the EU averaged series looks very similar, while the Malin Head result comes back to near zero at about 50 GW installed. In Figure 37 we can easily distinguish two regions: one for up to about 5% of the total electricity demand covered by wind energy, with a rather steep slope, and another domain for higher penetration. The steep slope in the beginning is 60%, meaning that the capacity credit for wind energy is more than 60% of the installed capacity. This value seems rather high, when the load factor of the wind time series is only 30%. Especially since the result is identical for the EU-Averaged series, which has an overall load factor of only 22%.

It had been pointed out before [130, 131] that the LOLP *Ansatz* used with a chronological model like the NGM is highly sensitive to single events. Hence, it is required to look at it in more detail on a day-by-day basis. The most reasonable place to look for a single event which could influence the common behaviour of all the time series is the day with the highest load and hence the highest demand on fossil fuel plant. The model time step with the highest fossil fuel usage occurs at hour 1628, seen in Figure 38. Here, we find the reason for this extraordinary number: the wind power at this point is 3918.4 kW, one of the highest values overall, or equivalent to 64.2% of the installed wind power capacity.

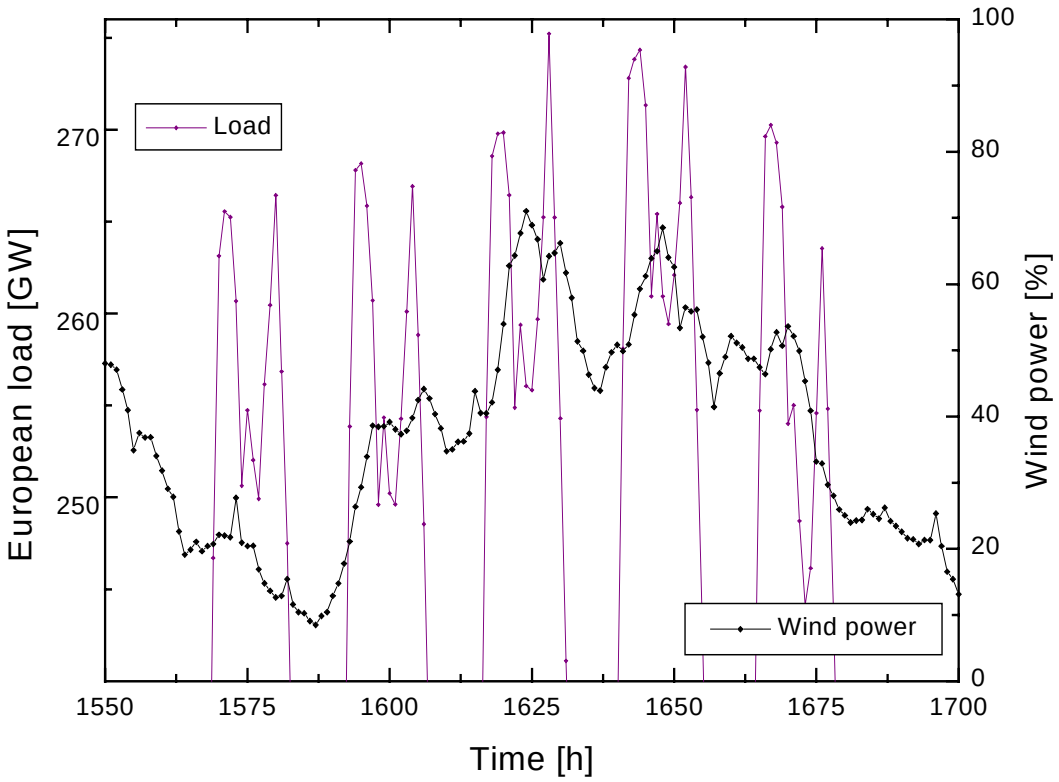

Figure 38: Load and wind power time series at the hour of the highest load.

So why does this value get so much smaller when installing over 15 GW of wind capacity? The answer, again, can be found in the details of the time series: at about 15 GW installed capacity, the wind power production is just enough to cover so much of the demand on hour 1628, that another hour takes over as the hour with the highest fossil fuel demand. This behaviour is illustrated in Table 4:

Table 4: Wind power and fossil fuel load on selected days. Days in bold are the days with the highest demand of fossil fuel plant.

Added Wind Capacity [GW]	Day	Wind Power [GWh]	Fossil output [GWh]
0	**69**	0	3995
2.745	**69**	36	3960
8.235	**69**	108	3887
13.725	**69**	179	3814
27.45	**69**	359	3635
82.35	62	359	3338
	69	1076	2936
	75	206	3544
137.25	**75**	343	3390

In this table we see that day 69 has been overtaken by day 75 as the day with the highest demand for fossil fuel electricity generation. On this day, the wind power output was down to only 11.9 % LF. This is the slope we found in Figure 37 for the higher installed capacity part.

A change was then introduced to get a more accurate representation of the nuclear plant strategy. In the previous graph, nuclear plant was supposed to have a higher efficiency in winter than in summer, since the Carnot cycle of the turbines has a better thermal efficiency if the lower temperature of the cycle is

lower. Later in the year, the nuclear plants are not delivering so much output. The reason for this set-up can be traced to the ancestry of the NGM, dating back to the 70ies. However, most modern nuclear plant run always at 100%, except for the summer months, when demand is low and the plants can be taken down for three weeks major maintenance and change of fuel rods. This was modelled with a two-month period in summer, when the nuclear plant contributed with only 70% of its usual amount. The results are rather different from before, not so much the financial result (which is essentially identical to Figure 36), but the unused fossil fuel capacity (Figure 39).

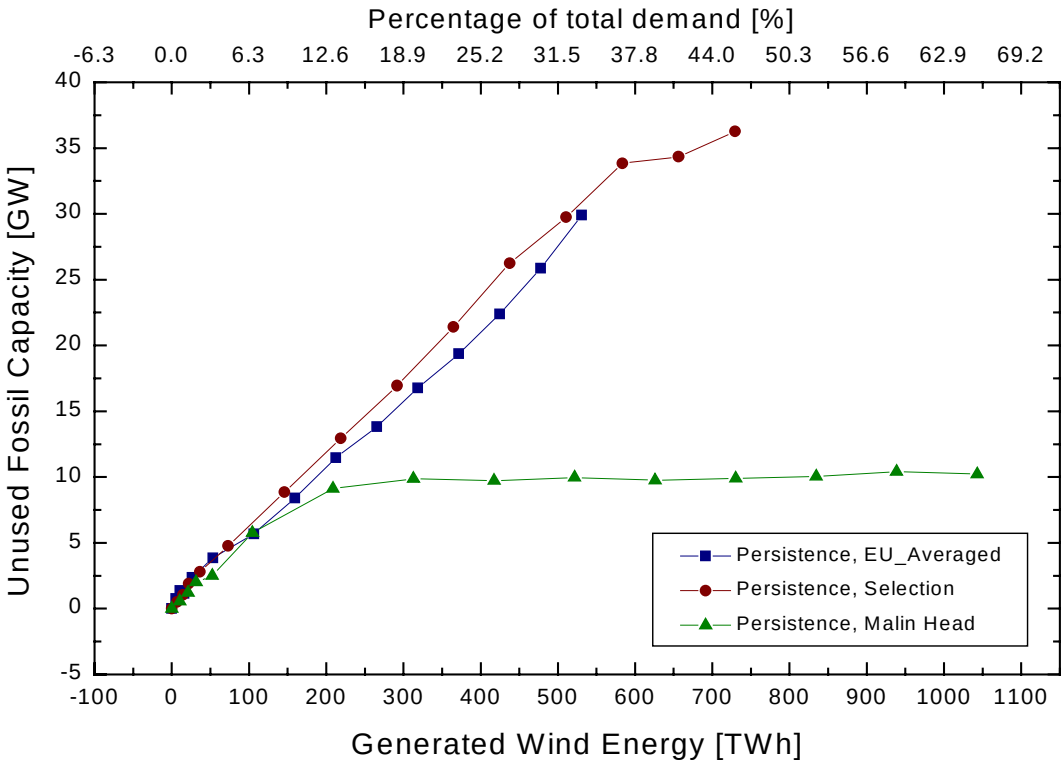

Figure 39: Unused fossil fuel capacity with the newer nuclear strategy.

Here, as before, the behaviour of the Malin Head time series is completely different from the others. Based on our previous experiences with single events, we expect to find this here, too – and we do. The following table gives a good account of the findings for day 75, where the load is rather high, while the amount of generated wind energy is near zero. Here it is very clear that the spreading out of wind farms is beneficial for the whole energy system, since the probability for wind energy production near zero is greatly reduced.

Another clear feature of Figure 39 is that the capacity credit for the other time series is much lower in the beginning, but stays there throughout the whole addition process of wind capacity. The unused fossil fuel capacity here is about 15.4% of installed wind power capacity for the selection time series, which corresponds again to the load factor at the day of the highest fossil demand. In this version, the determining event occurs at day 55, hour 1316, since here, the nuclear plant did not deliver as much power as in the previous set-up. Hence, the capacity credit is comparatively poor.

Table 5: Wind power and fossil fuel demand with the newer nuclear strategy.

Added Wind Cap. [GW]	Day	Load [GWh]	Wind Pow [GWh]	Fossil Gen [GWh]
27.45	**55**	5739	59	3657
	75	5619	9	3587
54.9	**55**	5739	118	3598
	75	5619	19	3578
82.35	55	5739	177	3539
	75	5619	28	3596

This result can also be plotted to show the replaced fossil fuel capacity as a percentage of the installed wind capacity. This is shown in Figure 40. The decrease of displaced capacity with installed capacity is obvious. The relative capacity effect gets rather low at very high penetrations. One feature of this curve that is particularly striking is the value above 100% for perfect forecasting of Malin Head. The reason for this is the discrete plant size of the displaced plants. If a previously only part-loaded plant is replaced, then the whole plant will count towards the displaced capacity. An additional effect is that adding wind energy to the grid, the spinning reserve requirement can add an extra margin to the actual production at the determining hour. Similarly, the optimisation can yield a lower spinning reserve figure than without wind. Actually, rather low wind energy generation at the hour of highest demand can lead to a higher overall spinning reserve requirement, therefore also values below zero can happen. We will see this in the next section.

There is a reason why the incremental capacity credit has to decrease with penetration. Conceptually, the capacity credit (or rather the unused fossil fuel capacity) is filling up states defined by the load minus

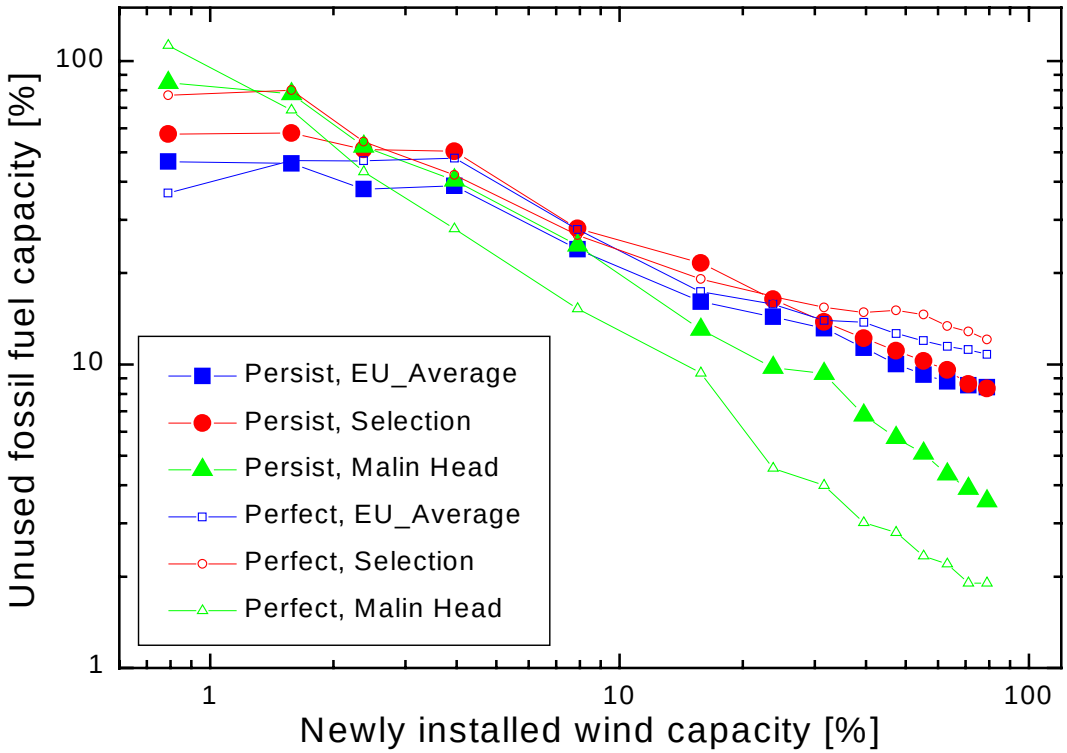

Figure 40: Displaced fossil fuel capacity as a percentage of installed wind generating capacity. Note the double log axes.

the base load generation. Base load generation is nuclear and hydropower combined. Base load is fixed in the model, therefore wind power can only replace thermal and other plant. It starts with the deepest state available (s_1, defined as *max(load-base_load_generation)*) and fills this state up. The rate at which this state is filled is determined by the load factor at that hour (lf_1). Therefore, with additional wind capacity the state is going to be filled in until the wind power capacity reaches the point where another state (s_2) takes over. Defining pen_1 as the first amount of newly installed capacity for which another state takes over, then this is given by

$$s_1 - lf_1 * pen_1 = s_2 - lf_2 * pen_1. \qquad (16)$$

This leads to

$$pen_1 = (s_1-s_2)/(lf_1-lf_2). \qquad (17)$$

If there is to be another state that takes over, then it has to have a lower load factor. *Ie*, demanding that pen_1 is positive, and taking from the definition that $s_1>s_2$, it follows that $lf_2<lf_1$. This argumentation continues for every following state and load factor accordingly.

This behaviour is exemplified in Figure 39. For the Selection time series, the first state has a relatively low load factor. Most other similarly deep states seem to have similar or higher load factors. Therefore, s_1 can retain its position all the way up to 60% additionally installed capacity (the point at 36% total demand covered), while then a state with even lower load factor takes over as the determining state, leading to a correspondingly lower slope. Actually, this picture is not quite right, when we look at Figure 40. Here we see that the first two values are flat, while only then another value takes the lead for a long span of installed capacity.

There are two interesting points that can be explained nicely with this state picture. First, the capacity credit for very small penetrations is determined by the value at the highest load event (actually, the highest *(load - base_load_generation)* event). In other words, if *pen=0*, then the relative displaced capacity is the load factor lf_1. This is to say that using a chronological model, the result of the capacity credit assessment depends on the load factor at the hour with the deepest state. On first glance, the wind speed and the highest load do not seem correlated. Therefore, the capacity credit of wind energy for small penetrations is more or less randomly distributed in the same way the wind power generation is distributed. However, as we will see later in chapter 8, the wind power generation is higher in winter, when also the load is highest.

There is not much literature on the correlation between wind speed and load. The correlation between temperature and load is well known and is used by utilities for load forecasting, but the correlation of temperature with the wind speed or an explicit correlation of wind speed with load has not received attention. Nevertheless, Nielsen and Madsen [132] analysed the correlation between wind speed, temperature and electricity demand using 20 years worth of hourly time series from Jutland/DK. They found a weak correlation between the wind speed and the electricity demand, which was much less pronounced than the temperature dependency. However, while in the early years (1974-1989) the correlation was slightly positive with wind speed, the higher amount of installed wind power in the grid reversed the slope of the dependency for the later years. A relation between slope and installed wind power could be found. Therefore, it can be concluded that the result for the earlier years is

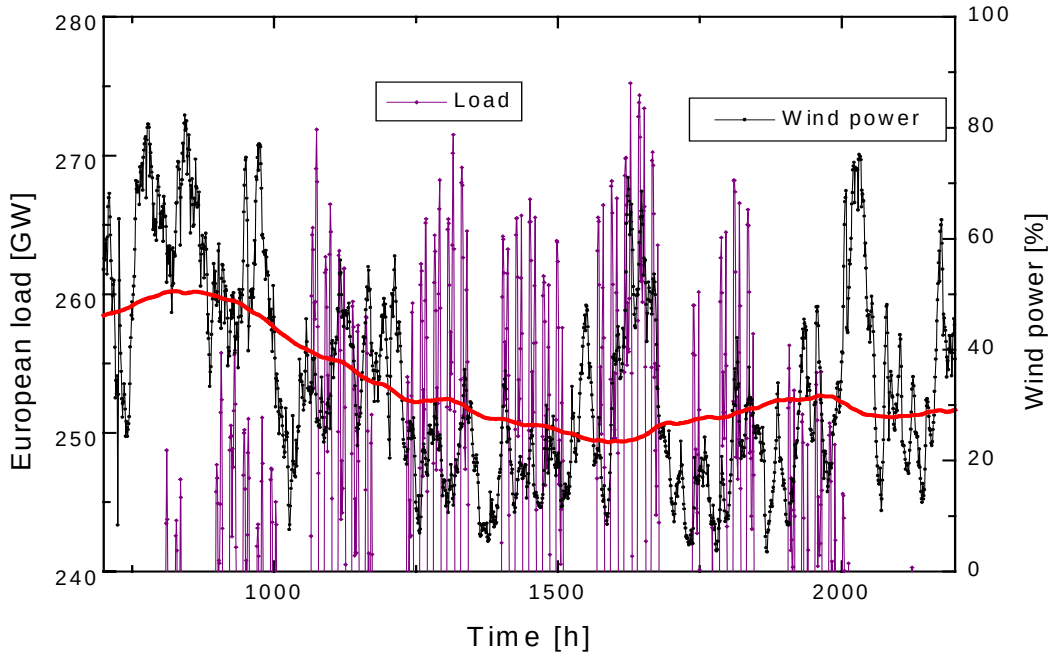

Figure 41: Kernel smoothed wind output and load. The graph contains the 7-week period with the highest demand. The kernel window was equivalent to one month (thick smooth line).

representative for the 'real' dependency, when not considering wind power generation at all. This also means that there is a slight positive correlation between peak electricity demand and wind speeds. This dependency was not just an artefact of the correlation of low temperatures and high wind speeds in winter, since the temperature dependency was modelled explicitly. This could be attributed to the 'chill factor' of wind, especially in a region where some houses use electricity for heating.

The other feature explained by the state picture is the behaviour for very large penetrations. Since the load factor (and therefore the marginal increase of the capacity credit) decreases monotonically, for very large penetrations it will bottom out at the lowest value. Let *pen* and therefore $i \to \infty$, then the relative displaced capacity goes towards the lowest load factor times a constant/*pen*. Therefore, the capacity credit for large penetrations decreases towards the lowest load factor in the wind power distribution. For small areas of wind energy generation, this will be zero power. However, as we have seen in Table 3, for all of Europe the generation stays (slightly) above zero.

Since we now have a random distribution of results, which one is the one that really counts? To assess this, we have to look at the load data first, to find which periods are likely to contain the important events. The overall highest demand is found in a seven-week period in winter (Figure 41), ranging from the second week in January to the third week in February. All the highest events we already encountered are contained in this period. The variability of the load over many years is quite small – so this seven-week period should be more or less the same over all years, barring extremely cold winter periods outside of these seven weeks. Therefore, it is important to learn about the wind in especially this time. This is also shown in Figure 41. Here, the wind is averaged with a moving average over a whole month, to give a reasonable amount representative for the period. We see that in this example the wind is quite strong during the highest load period, giving a load factor of roughly 30% (about 1800 kW production).

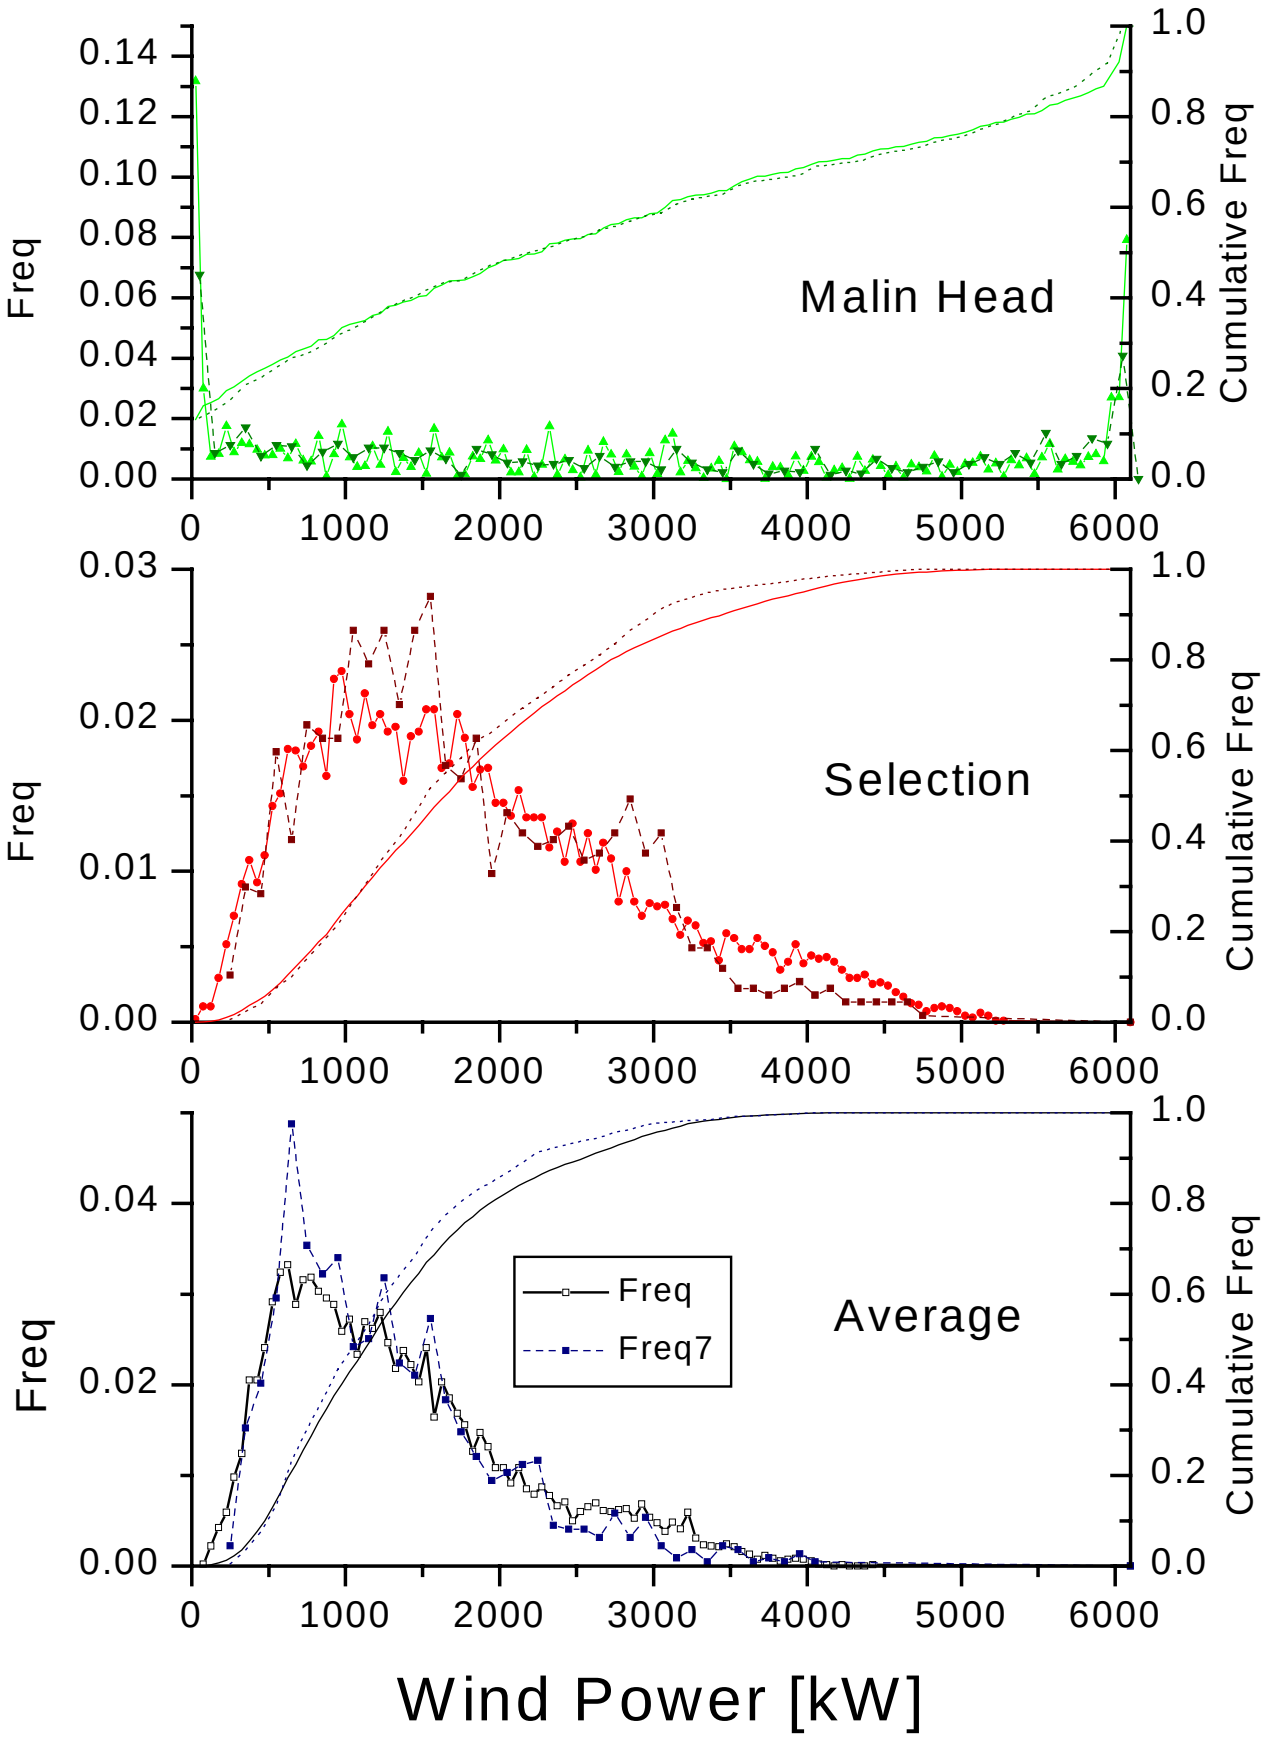

Figure 42: Frequency distribution and cumulated frequency for the three data sets. The dashed lines refer to the 7-week period, while the full lines are the distribution of the whole year.

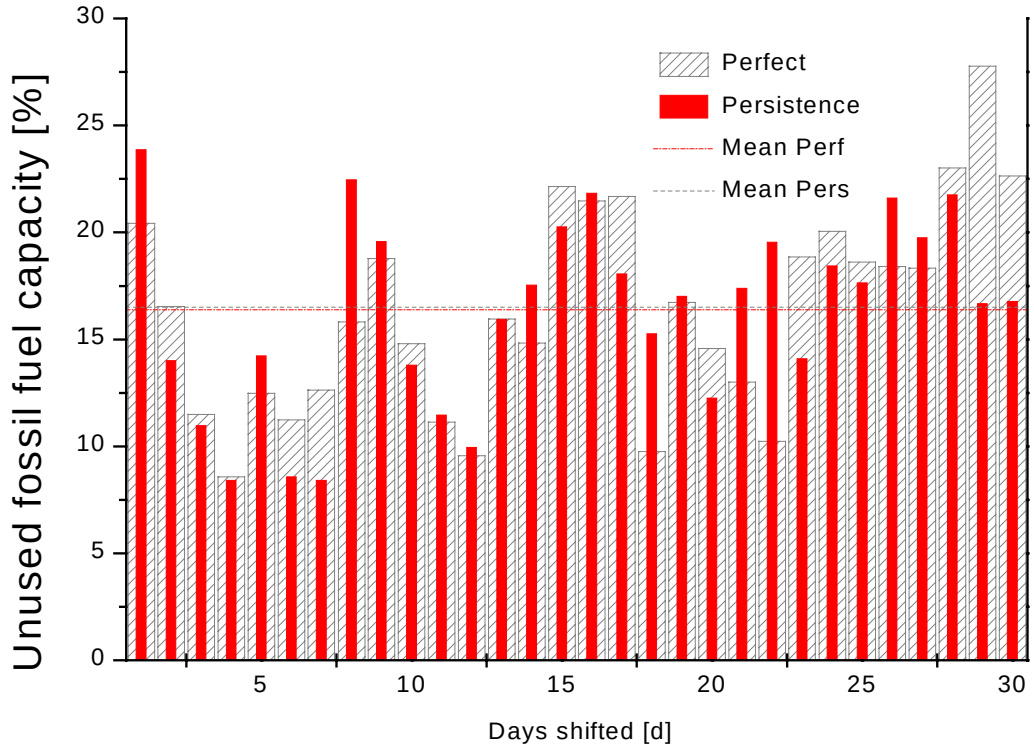

Figure 43: Variational analysis for the continuous time series. The Average time series was used.

In many cases, only a yearly frequency distribution of the wind will be available. Therefore, it is important to compare the yearly distribution with the distribution of the wind in the seven-week period pointed out here. In Figure 42, a frequency distribution analysis has been done for the three wind power time series, once looking at the whole year and once only at the 7-week high load period. As we might expect, the case of Malin Head is quite different from the other two averaged series. Both zero output and full output occur quite frequently, while the middle ground is fairly equally distributed. This points to the middle ground being run through on the way between full and zero generation without any particular generation being favoured. However, looking at the Selection and Average series it is somewhat surprising that the generation during the 7-week period in winter is slightly lower than the production during the whole year. This can best be seen when comparing the dotted line with the full line in both graphs. These represent the cumulative frequency. A faster rise of this line means therefore, that the distribution is shifted to lower values. In most cases, the production during the winter months should be higher than during the whole year – therefore, this result is surprising. To which extent the difference is typical or just an extreme case of the one year studied, will be shown in chapter 8.

7.6 Variational analysis

To overcome the single event problem encountered in the previous chapter, in this section the basis of our analysis should be broadened. This can be done by shifting the wind time series relative to the load time series. Since the load time series exhibits a diurnal variation, it was decided to only shift the series by whole days. 13 months of wind data were available, while only 12 are needed for one run of the NGM. This meant that a shift of up to 31 days could be done without destroying the continuity of the series. The analysis shown in Figure 43 was done for the target wind capacity of 40 GW laid out by the

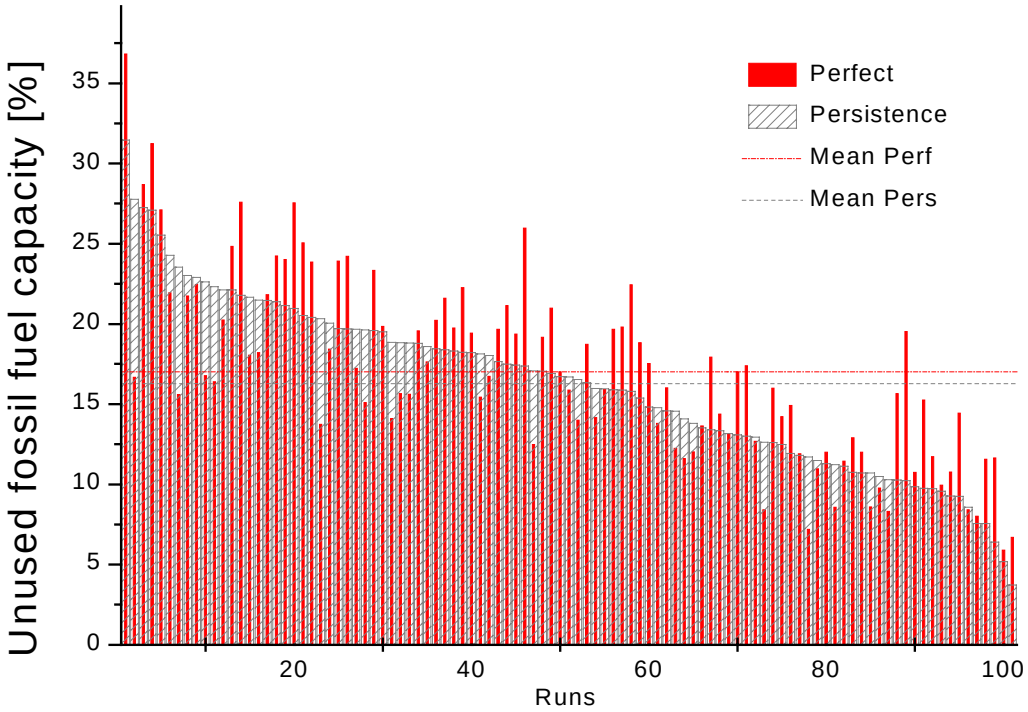

Figure 44: Unused fossil fuel capacity for a variation of ±50 days for the Average time series.

EU commission for 2010 in the white book on renewable energy [133]. In the power plant mix used here, this corresponds to 11.5% additionally installed wind capacity, or 10.3% penetration.

The variation found in these 30 days is already quite significant. However, the mean unused fossil fuel capacity is for both persistence and perfect forecasting 16.5% of the installed wind power capacity. The minimum encountered in this period is 8.5%. The corresponding numbers for the Selection series are generally higher, but also show larger variation.

If we drop the requirement for an undisturbed time series and allow a wraparound of our data, we can also shift the data for more (padding non-existing data in the winter with data from the other end). This was done here for ±50 days, since using even more days, the synchronicity of the seasons would have been neglected. Since the variability is not quite well to see in a chronological graph like the one in Figure 43, for Figure 44 the data has been ordered by decreasing unused capacity using persistence forecasting. The highest and lowest values are easy to find here. The minimum is at 3.8% and 5.9% for persistence and perfect forecasting, respectively. The difference between perfect and persistence forecasting stems here from the different spinning reserve requirements - perfect forecasting facilitates on average a lower spinning reserve requirement. Since the scheduling of power plants has to take the spinning reserve into account, the unused fossil fuel capacity is (in the framework of the states introduced in the previous chapter) calculated against slightly different states. Therefore, perfect forecasting can slightly improve the result here. For perfect forecasting, the mean is at 17%, while for persistence forecasting the average unused capacity is 16.3% of the installed wind capacity. However, for the Selection series perfect forecasting performs worse on average than persistence forecasting (19.8% and 20.5%, respectively). The difference seems random, and might be explained with the different sizes of power plants displaced in the different states.

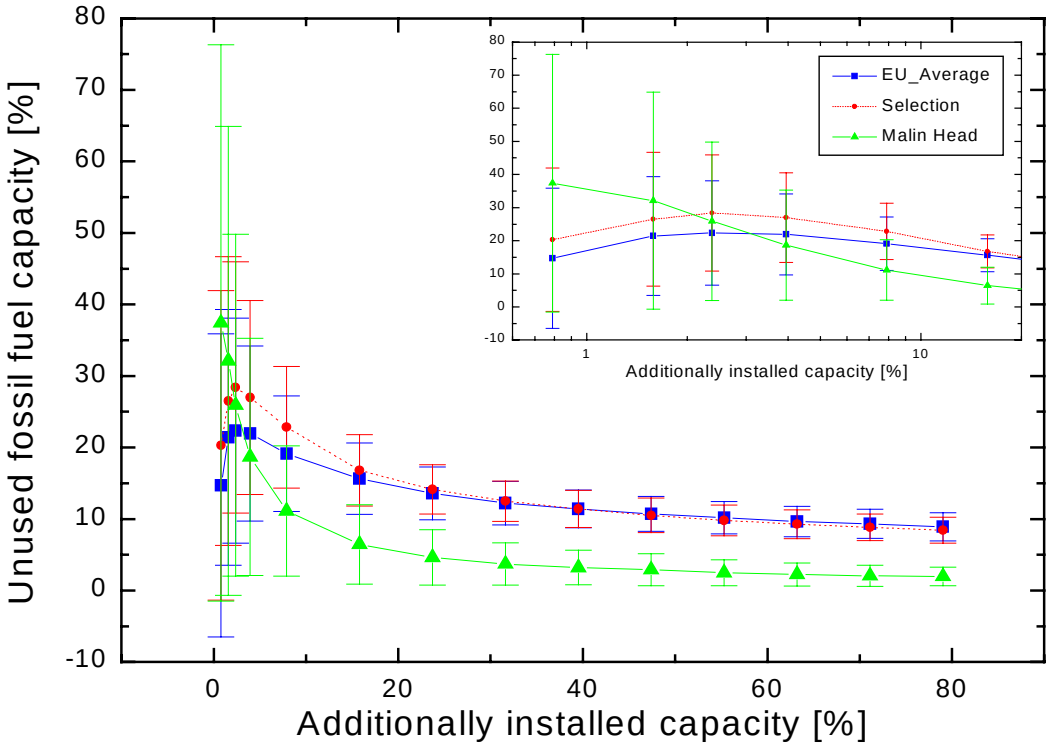

Figure 45: Variation of the unused fossil fuel capacity with increased penetration. In the inset, the same graph is shown for small penetrations. Note that the scale in the inset is logarithmic, for spreading out the calculations done for small penetrations.

For Figure 45, all the instances of relative displaced capacity have been calculated analogous to Figure 40, using the shifting process and perfect forecasting. The average displaced capacity is plotted together with the standard deviation of the instances. The explanations given by the state picture are followed closely for the case of Malin Head: for low penetrations, the mean displaced capacity is close to the average load factor, while the scatter is very large. For high penetration instead, the mean displaced capacity decreases with a low scatter to near zero. For the two other graphs, the behaviour at high penetration is also explained by the state picture: the Average time series has a higher minimum load factor than the Selection series (see Table 3). Therefore, at the highest penetrations analysed, the relative displaced capacity of the Average series is higher than the one of the Selection series. However, the state picture does not explain the increase in relative displaced capacity. The existence of a maximum has also been seen in Spain [115] (see p.59). Here, the explanation with the discrete power plant size is likely, too: the unused fossil fuel capacity for 0.8% newly installed capacity is for the Selection series at 550 MW, a size where one or two typical coal fired power stations are being replaced. In most cases, it replaces only one, while reducing the load factor of the other. That other one is then still counted as needed.

7.7 Conclusions

A capacity credit assessment of wind energy using a chronological model is difficult, since single events tend to dominate the behaviour of the result. Note that the financial benefits from saved fossil fuel are largely unaffected. One way to reduce the insecurity of the displaced capacity is to do a variational analysis, shifting the wind power time series against the load data. Thereby, the variation on a time scale

of days is accounted for. The single events most likely to influence the result are in a seven-week period in January and February, when the load is highest. The wind during that period of the year is therefore the most important one for the calculation of the capacity credit. The distribution of the wind during the high-load period is determining the behaviour of the displaced capacity, for small as well as for high penetrations. Since the wind speeds during a period are important, no single value for a capacity credit can be given. For large penetrations, it decreases towards a value depending mainly on the minimum load factor. For the European Average wind analysed here, the unused fossil fuel capacity is about 9% of the installed wind capacity at 45% penetration. For small penetrations, the relative displaced capacity will be on average close to the average load factor during the important period. It will scale with the load factor at the time of the highest demand. This is higher in winter, when also the demand is higher. There is a positive correlation between wind speed and demand. Therefore, the determining load factor is higher than the average yearly load factor.

Another result from the analysis is that perfect forecasting does not necessarily lead to a better capacity credit than persistence forecasting. On the other hand, perfect forecasting allows more wind energy to be used in the grid. This behaviour is especially pronounced with very variable wind power generation.

For the relatively smooth average production in Europe, 20% of the total demand can be covered, while discarding 10% of the generated wind energy. This percentage of wind energy could probably be used up by effects not modelled in the National Grid Model, such as hydropower reservoirs. At this stage, wind energy would save nearly 60% of the total fossil fuel cost of electricity generation in Europe, worth close to 7 G€. This would give an estimate of 2.2€c/kWh as the worth of wind energy in fuel saver mode.

8. The long term - Reanalysis Data

In the course of the last chapter, the importance of the load factor during the period of the highest load for a capacity credit assessment has been stressed. In Figure 42 we saw that the winter generation in 1991 was slightly lower than the yearly average. The question is: is this a general feature of the European wind power generation, or is this a special case in 1990/1991? Ideally, one would use a measured time series of 10 years or more at the sites in question. The lack of detailed data for more than one year can be overcome with the use of reanalysis data. See section 3.8 for an explanation. Reanalysis data was available for 34 years, therefore an assertion can be made regarding the long term effects on the previous analysis.

In Figure 46, the meteorological data time series is compared to data derived from reanalysis. The meteorological time series is the Average series from before, while the reanalysis data was similarly generated. The cross-correlation of the two data sets is 0.87. If used only for sites with more than 2000 FLH (27 in total with reanalysis), the correlation drops to about 0.72. This is about the same correlation as two wind time series from 200 km apart. The poor spatial resolution of the reanalysis data grid might be able to explain this, where the average distance between the grid point and the meteorological station is in that range. Another factor is that numerical weather models cannot perfectly predict the wind at a given station.

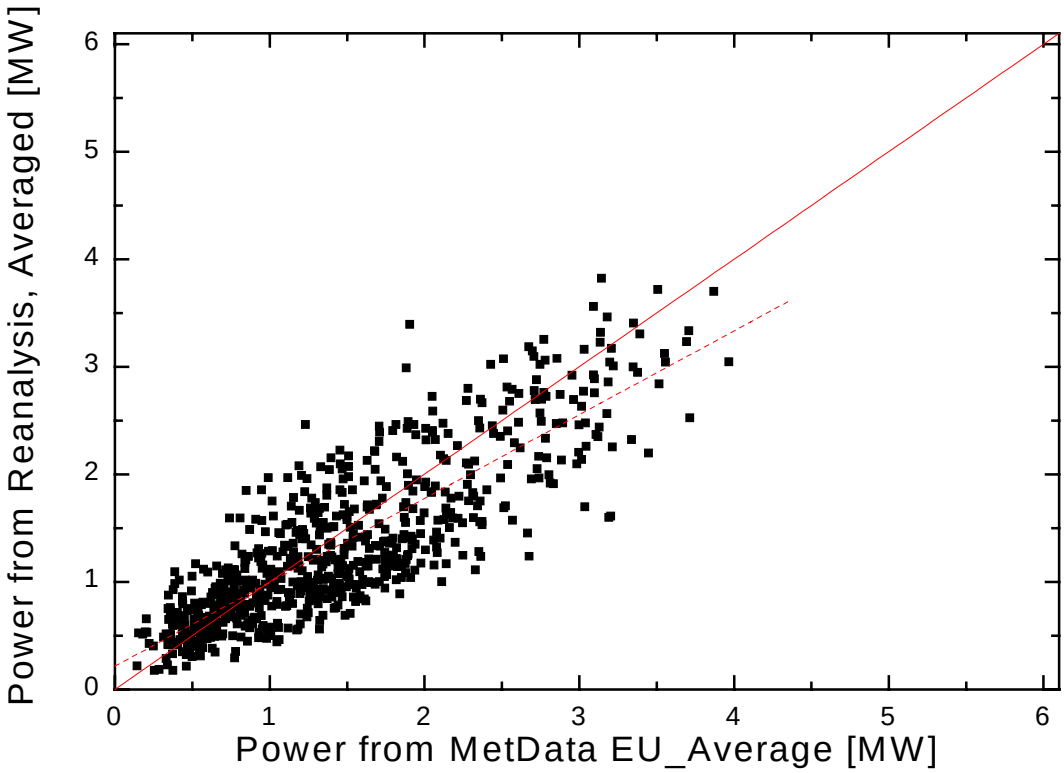

Figure 46: Comparison of meteorological data and reanalysis time series. Power from Met data refers to the averaged time series derived from meteorological data, while Power from Reanalysis refers to the time series synthesised from reanalysis data. The period is from Dec 1st, 1990 to Nov 30th, 1991. The full line is equality and just to guide the eye, while the dashed line is a least squares fit through the data.

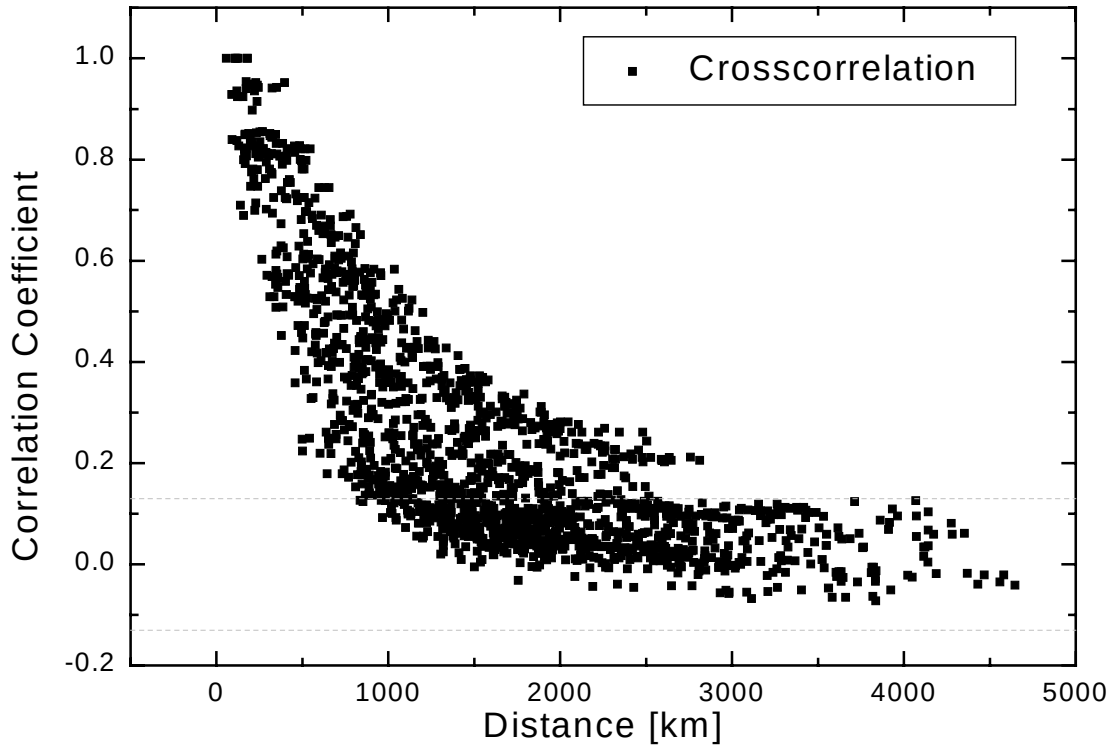

Figure 47: Correlation coefficient for the Reanalysis time series.

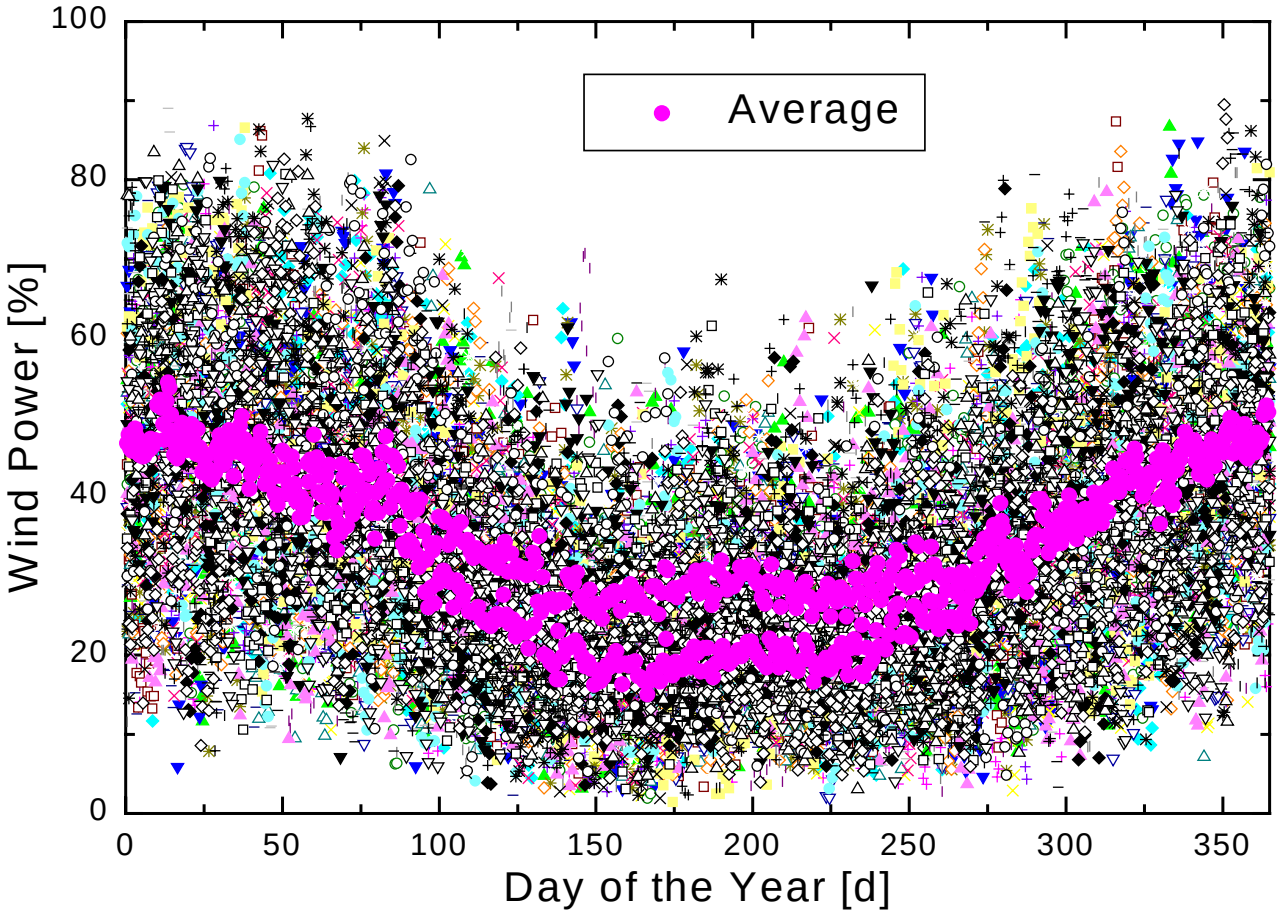

Figure 48: European generation according to reanalysis data. Every single time series corresponds to one year, while the Average is the average at every time step of the 34 years.

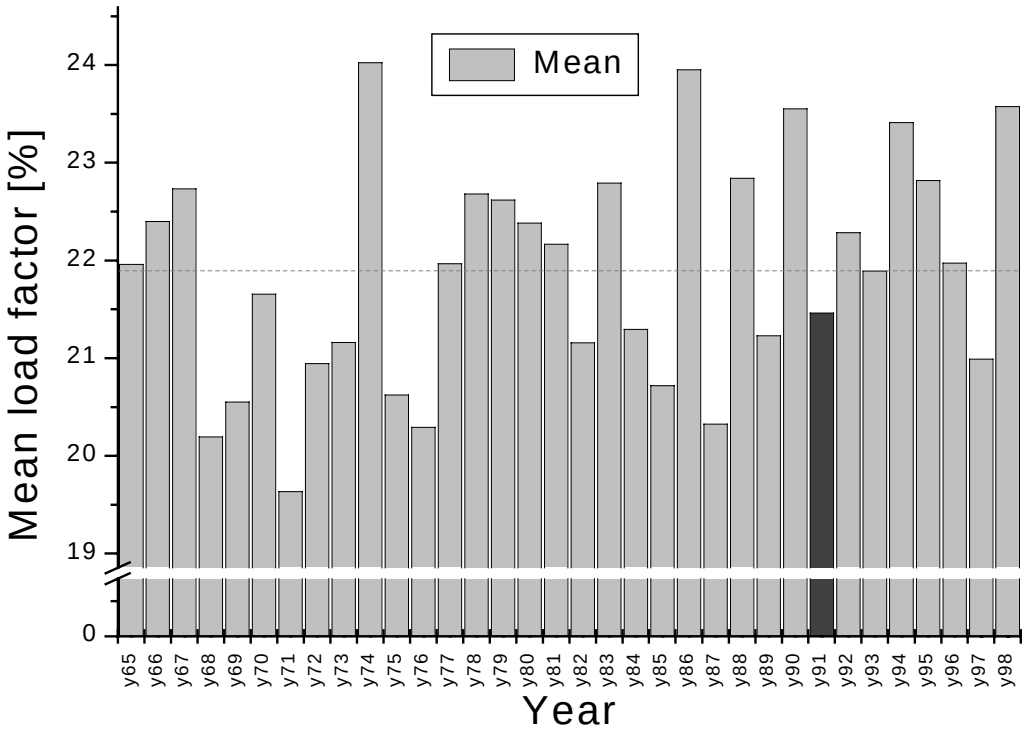

Figure 49: The 34 yearly mean wind power productions from 10m reanalysis data and the long-term mean (dashed line) for the Average series.

In Figure 47, the cross-correlation coefficient of the reanalysis time series is shown. An exponential fit to the data yields a decay parameter of 875 km, which is slightly more that in Figure 28. The reason could be that every data point from reanalysis is a wind characteristic for a large area (ca. 200x200km). Using a mean wind of 8 m/s, this corresponds to a 5-hour displacement. Therefore, the wind is akin to a 5-hour mean. The values of 1 are from the same time series. Since the nearest grid point was used, in four cases this meant using the same grid point from the reanalysis grid for two stations close to each other. It can be concluded that the reanalysis series is approximating reality quite well, although it might underestimate the variations slightly.

A typical feature of the wind energy production is that in summer wind energy production is much lower than in winter. The different yearly time series are rolled out in Figure 49. Every point shown here is one realisation of a wind energy production, as averaged over all Europe, in 34 years. That means, at every of the 730 time steps available from reanalysis (365 days, two values a day), 34 points are scattered along the y-axis. More interesting than this cloud of points is the empty area surrounding it. Here, in no case during the 34 years analysed, production occurred at this level. This is all the area above 90% generation, and above 70% during the summer months, but also the area below 10% for the winter months, where only very few cases of low wind are seen. As mentioned before, reanalysis probably smoothes reality, and therefore the lack of data there should be taken *cum grano salis*. However, it is a quite strong indication that values in these areas are rather unlikely.

It can also be seen from this plot that on average, there is a daily variation for the summer months. This is thermally driven, and since the higher wind is at midday, this coincidence with the cooling need in Southern Europe with wind is fortunate.

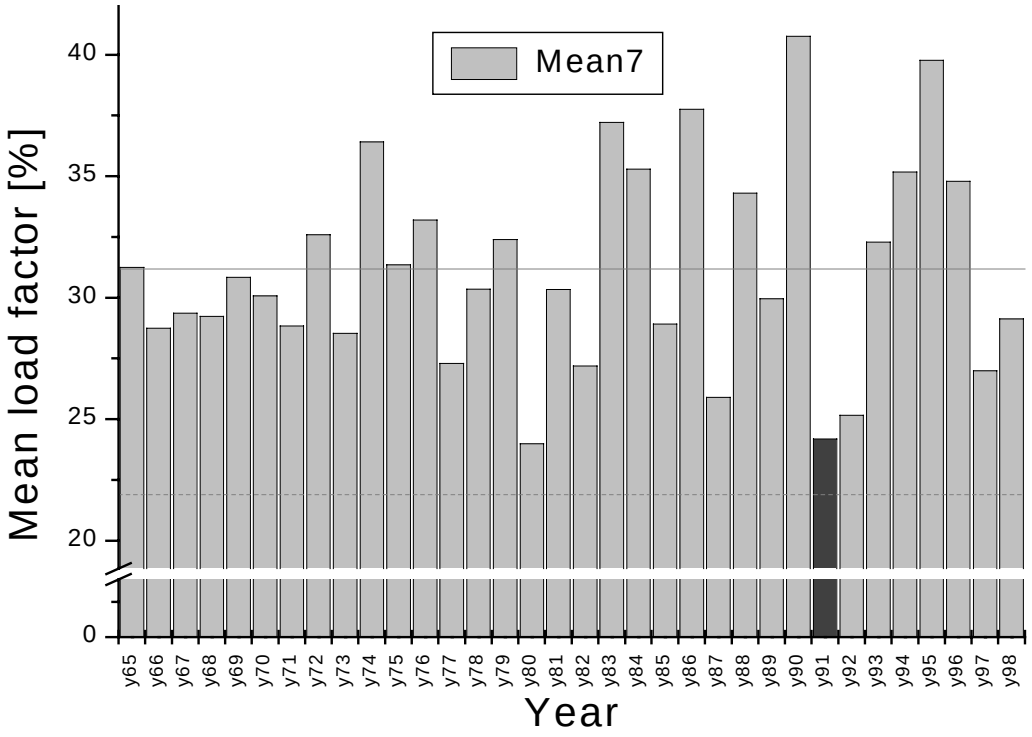

Figure 50: Mean load factor for the Average series, averaged over the 7-week period of the highest load. The full line shows the average over all columns, while the dashed line corresponds to the overall mean power production (identical with the line from Figure 49). Please note the different y-scales in Figure 49 and here.

Let us now come to the analysis. In Figure 49 we see that the wind power has quite some spread over the years, and 1991 was a just-below-average year. This means that the analysis done here probably slightly underestimates the value of wind energy in comparison to the long-term mean. We have also seen that the important period for a capacity assessment is the 7-week winter period. It could also very well be that the relation of the wind speed during the winter period to the yearly average wind speed is not uniform over the years - in fact, this is to be expected.

Figure 50 shows the mean power production for the 7-week period of the highest load. Our reference year 1991 is actually one of the worst producing winter periods in the whole period. This is even better accented in Figure 51: Here we see that in the mean the average power production during the 7-week period in question is 35% higher than during the whole year, while in 1991 this ratio was only 11.9% better. This is another indication that 1991 was one of the worst years to use for a capacity assessment of wind energy. In other words, the results presented in the previous chapter should be treated as a worst case scenario.

The final issue is to do an analysis of the frequency distribution analogous to the one in Figure 42 for the long-term average. Here, all power generation values of the 34, 7-week periods are used to get the probability distribution. In Figure 52, which was calculated using the Selection time series, the most interesting is the again cumulative line: here we find that 95% of all wind power during the highest load period are above 1180 kW, which translates to a capacity factor of 19.3% of the installed capacity. The mean is at 2805 kW, or 46% respectively. The corresponding numbers for the Average series are 11.8% and 23.6%, respectively.

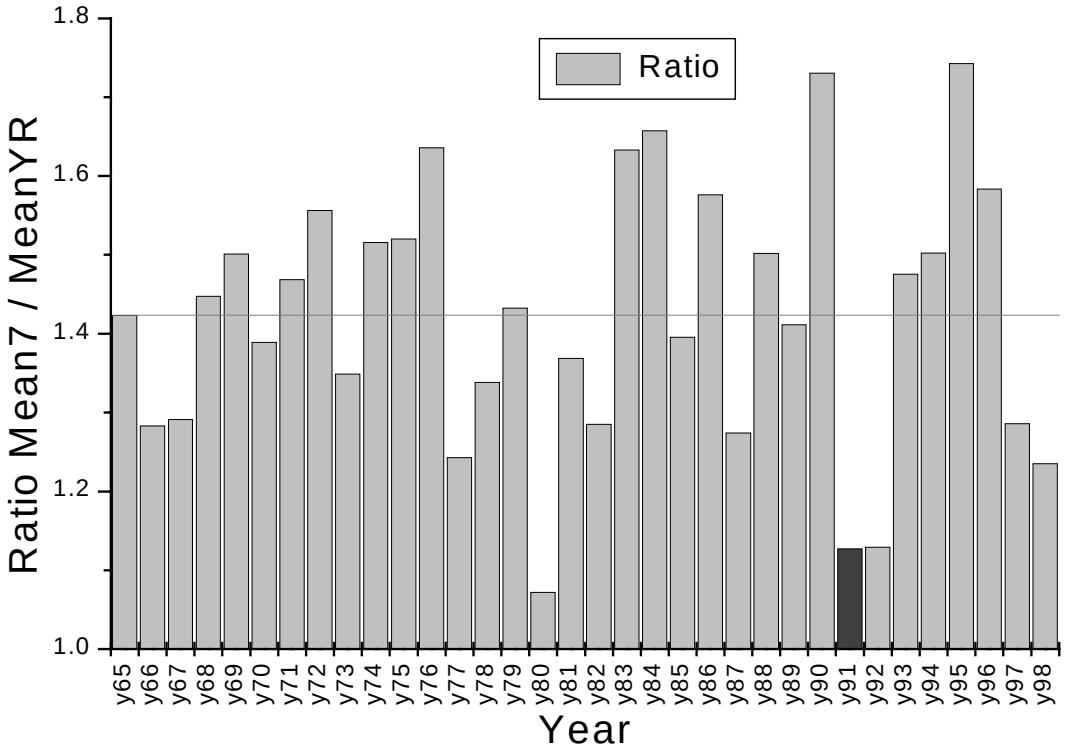

Figure 51: Ratio between the mean power production during the 7-week period of the highest load and the average production during the whole year. The full line is the average ratio.

The important point for a capacity credit assessment is the difference in mean power output averaged over the whole year (a number which can fairly easily be obtained from *eg* the internet) and the mean power output in the seven-week period in question: the overall mean production is 2074 kW (or 34% of the installed peak power), while the mean production during the highest load is 2805 kW (46%). Even though the actual numbers are dependent on the choice of locations and the power curve, the general fact remains that the mean load factor to be used for a capacity credit estimate of wind energy is more than a third higher than the load factor averaged over a year.

Concluding can be said that the reanalysis method works. It could be used to set the results of the previous chapter into a long-term context. The year 1991 used in detail in the previous chapter can be regarded as a worst case, since the generation in the 7-week period of the highest load is one of the lowest in all the 34 years analysed. Therefore, a capacity credit assessment based on this year will give a worse result than with most other years. The general conclusion for a capacity credit assessment is that the wind power generation during the high-load period is on average 1/3 higher than the yearly mean generation.

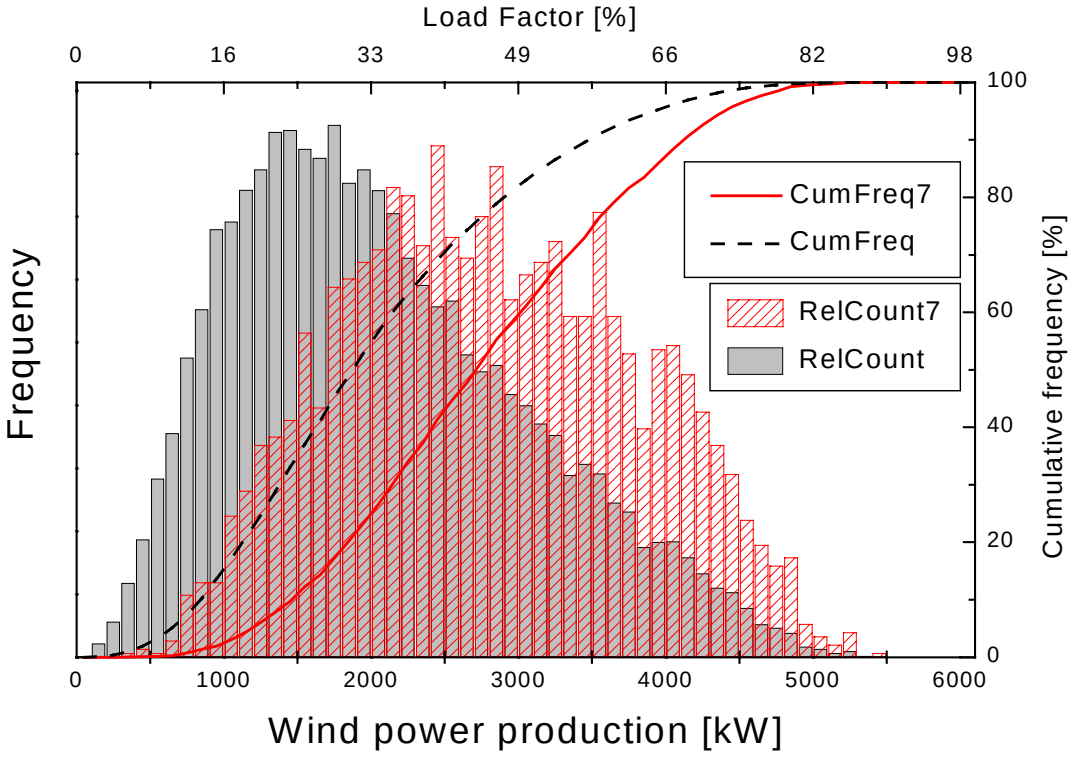

Figure 52: Frequency of occurrences for the wind power output of the wind power time series during the seven weeks of highest load of all 34 years. The grey columns and the dashed line refer to the distribution during the whole year.

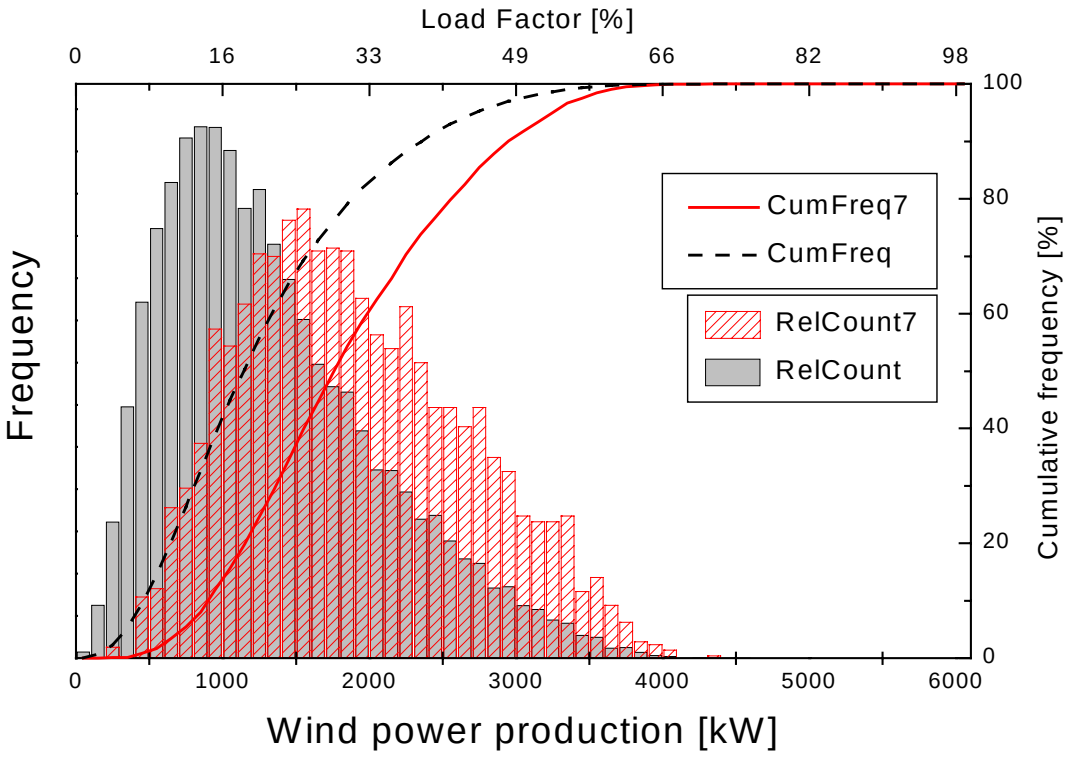

Figure 53: Same as Figure 52, this time for the Average series.

9. Summary/Conclusions

The aim of this work was to improve the integrability of wind power into the European grid, tackling the task from a technical side and from a conceptual side. The technical issue was to try to improve the quality of short-term predictions of wind power, to help utilities integrate more of this variable resource more easily. The other part was trying to change the impressions often found about wind energy: that it is too variable a resource to be relied upon. Distributing the resource over all of Europe significantly decreases the variability of the resource, thereby facilitating savings in investments in the power plant mix.

The Risø model of using HIRLAM, WAsP and PARK for short-term forecasting of wind energy is state-of-the-art. However, occasionally error sources are introduced that are not modelled by this set-up. Model Output Statistics can be used successfully to improve the output. The improvements possible are mainly dependent on the farm, and not on the forecast horizon. This points to the HIRLAM wind as the main culprit for errors. The Kalman Filter / Extended Kalman Filter can be used as a recursive module, giving improvements over model implementations with fixed parameters. In no case did the KF perform worse than fixed parameter models. When using it, there is a significant difference in performance depending on the quality function - usually there is a trade-off between optimising for the Root Mean Square error and the Mean Absolute Error. The benefit of the Kalman Filter is most pronounced for situations where the seasonal variation built into HIRLAM is not the same as the one on the site, or for changes in HIRLAM, where a fixed filter would have to be set up again. However, since the filter has to be rather stiff in order to not be irrecoverably taken off course, the possible improvements over fixed parameter models are limited.

The assessment of the capacity credit of wind energy by means of a chronological scheduling model is very sensitive to single events. The scheduling model can be used nonetheless to assess the fossil fuel savings due to wind energy. It also can help in estimating the benefits of distributed generation or various forecasting options. An important result is that the averaging of many wind time series from all over the EU is beneficial for the reliability of the wind resource, since wind power is available somewhere in Europe at all times. Using data from 60 meteorological stations all over Europe, the resulting wind generation time series was fed into the scheduling model. Wind energy can contribute more than 20% of the European electricity demand, taking a loss of 10% of the generated wind energy into consideration. Under the assumptions of the scheduling model, this leads to savings of 60% of the fossil fuel cost, worth about 7 G€. The installed wind power capacity replaces at this stage somewhat more than 10% of the installed capacity. Using a 34-year time series derived from reanalysis data this could be shown to be a worst case scenario, since the analysed year was one of the worst in terms of average load factors during the high-load period. Generally speaking, the average load factor relevant for a capacity credit assessment is more than a third higher than the yearly load factor of the time series.

Here also some ideas for future work: The benefits of distributed generation could be analysed with more data, especially load and wind data. Other countries could be analysed, especially the hydro storage possibilities of Norway, and the transmission issues could be looked at in more detail. The hydro storage handling of the National Grid Model is also due for an improvement. Currently, the new wind energy capacity is built in addition to the existing power plant mix, which is thought of as stable. However, due to the time needed to build up to the high penetration of wind energy assumed, the load as well as the power plant mix would probably change during that period. One could try to make educated guesses or scenarios on future developments in the energy business, and introduce a timeline with different reference scenarios along the way.

An analysis of the ideal data window for the calculation of the mean in the New Reference Model could yield interesting insights. Maybe it is possible to correlate the ideal window with the typical time scales of the wind speed variations at the particular site.

In general, it has been shown that wind energy can take its role in the future energy supply of Europe. In order to fully benefit from the distribution of wind energy generation, a common framework for market access has to exist in the whole EU. It must be possible for wind energy from everywhere to be consumed everywhere. This basic assumption has been made for this study. The benefits of distributed generation in Europe are not going to be realised if wind energy sales are restricted to the golden or not so golden cages built by the member states.

Acknowledgements

First of all, the institutional donors: I would like to thank the EdF for donating a load time series to the project, and I do thank the EU Commission for the Marie-Curie-Fellowship JOR3-CT97-5004, which liberated me (in conjunction with the European welfare state model and Risø's generosity) from economic worries. My special thanks in the EU Commission goes to Barry Robertson, who had the idea of organising workshops for the grant holders, to get to know each other and exchange ideas about work, life and everything.

Another set of thanks goes to OSTI.gov and the Virtual Library of Energy Science and Technology. Many of the ideas and clarifications I got were from papers effortlessly retrieved from their website. I would wish for an as easy access to European research, and the general attitude that research that once was paid for by the state should be in the public domain. Hopefully, CORDIS will get there one day…

In the same breath, I would like to thank the folks of the reanalysis project, for handing out their work to everyone with a web connection. Another thanks goes to the Scientific American, which made one special edition on the energy problems of humanity; this got me hooked thoroughly…

Closer to home, I would like to thank everyone who came with input - most of all Simon and Alfred, who always were willing to discuss with me. The same category applies to Igor, who is more than anyone else (perhaps excluding myself - perhaps) responsible for the form and coherence of this work. In the same bracket I would like to thank the teams at Risø, RAL and Oldenburg for a good working climate, scientific inspiration and the chance to experience foreign cultures.

Lars deserves a large amount of thanks for giving me the chance of at least starting a PhD (even though partly due to confusing my lack of oversight with persistence), and for guiding me during all these years. I would also like to thank Prof. Jürgen Parisi for accepting this work, even though it came from a complete outsider.

At last, I would like to thank the Gregorian Choir of St. Ansgar for (divine?) inspiration to the state model of the capacity credit. I am sure, I looked strange, sitting there hacking on a Palm Pilot, with everyone else celebrating Easter…

Appendix

Formulae describing the Wind Speed

The wind speed distributions at any given site is usually corresponding to a Weibull distribution, given by

$$f(u) = \frac{k}{A}\left(\frac{u}{A}\right)^{k-1} \exp\left(-\left(\frac{u}{A}\right)^{k}\right) \tag{18}$$

where *f(u)* is the frequency of occurrence of wind speed *u*. The scale parameter *A* is connected to the magnitude of the wind speed, while the shape parameter k has two special cases: for *k=1*, the Weibull distribution degenerates to a exponential distribution, while for *k=2* the function is called a Rayleigh distribution. The Weibull distribution is also the basis for the European Wind Atlas methodology [12] and the derived WAsP model [11]. Typical values for (especially Northern) Europe are 2-9 for *A*, and around 2 for *k*. The more the wind at the given site is driven by large-scale phenomena, the more it will generally adhere to the Weibull statistics. Typical reasons not to adhere are local thermally driven winds during daytime, or severe orographic effects like channelling in a valley.

The wind speed profile with height is determined by the friction velocity u_*. At large height a.g.l., the wind is not changing with height. This wind is called the geostrophic wind. It can be transformed to the surface using the geostrophic drag law, given as

$$G = \frac{u_*}{\kappa}\sqrt{\left[\ln\left(\frac{u_*}{f z_0}\right) - A_G\right]^2 + {B_G}^2} \tag{19}$$

where G is the geostrophic wind, f is the Coriolis parameter, and z_0 is the aerodynamic roughness length. κ is the Von Kármán constant, usually given as 0.4 [44, 50], though Bergmann [134] could show theoretically that the value should be *1/e* (≈0.37). A_G and B_G are constants, set to 1.8 and 4.5, respectively.

At low heights, the logarithmic wind profile uses the friction velocity derived from the geostrophic drag law to yield the height profile in the surface boundary layer of the atmosphere. It is given as

$$u(z) = \frac{u_*}{\kappa}\ln\left(\frac{z}{z_0}\right) \tag{20}$$

where *u(z)* is the wind speed at height *z*. Both equations are in their neutral form.

For a thorough primer on the meteorological concepts encountered, and how these apply to the problems at hand, I recommend the PhD-thesis of Kai Mönnich [50].

Statistical Formulae

In this thesis, mainly four error criteria are discussed: the Mean Error, the Mean Absolute Error, the Root Mean Square Error and the Skill Score. Assessed are the differences between the measured production P_{meas} and the forecasted production P_{fc} of a time series of length *N*.

The Mean Error ME is given as:

$$ME = \frac{1}{N}\sum_{1}^{N}\left(P_{meas} - P_{fc}\right) \tag{21}$$

The Mean Absolute Error MAE calculates as:

$$MAE = \frac{1}{N}\sum_{1}^{N}|\ P_{meas} - P_{fc}\ | \tag{22}$$

The Mean Square Error MSE is simply:

$$MSE = \frac{1}{N}\sum_{1}^{N}\left(P_{meas} - P_{fc}\right)^2 \tag{23}$$

Therefore, the Root Mean Square Error RMSE is defined as:

$$RMSE = \sqrt{MSE} \tag{24}$$

As a measure for the variability of a time series, the Variance VAR comes in handy:

$$VAR = \frac{1}{N-1}\sum_{1}^{N}\left(P_{meas} - P_{fc} - ME\right)^2 \tag{25}$$

From this, the Standard Deviation σ is derived:

$$\sigma = \sqrt{VAR} \tag{26}$$

To evaluate the predictions from different sources, the multiple correlation coefficient ρ given by

$$\rho = \frac{VAR(w_{t+k}) - MSE_k}{VAR(w_{t+k})} \tag{27}$$

is used. Since no absolute values are present in this skill score, it can be used to compare different sites and predictions. *VAR* is the variance of the (centred) time series of the observations and MSE_k is the mean square error of the predictions *k* hours ahead. The interpretation of this coefficient is that it measures how much of the total variation in the observations is explained by the predictions, i.e. 1 means that the predictions are perfect and zero means that predictions are useless.

This coefficient is actually identical with the Skill Score, usually expressed as

$$S = \frac{MSE_{ref} - MSE_{Model}}{MSE_{ref}} \tag{28}$$

MSE refers to the Mean Square Error of the reference model (the one to compete against) and the model to evaluate.

The correlation function of two time series p_t and q_t is as follows:

$$a_k = \frac{1}{N}\sum_{t=1}^{N-k}\hat{p}_t\hat{q}_{t+k}\ \Big/\ \sigma_p\sigma_q \tag{29}$$

with

$$\hat{p}_t = p_t - \mu_p \text{ and } \hat{q}_t = q_t - \mu_q\,. \tag{30}$$

$\mu_{p/q}$ is the mean of the corresponding time series, $\sigma_{p/q}$ is their standard deviation. *k* refers to the time lag between the two series. For the autocorrelation function, set $q_t = p_t$.

A value of 1 means that the time series are completely correlated, while a value of 0 means that the data is completely uncorrelated.

References

1 Available from *ftp://ftp.ripe.net/iso3166-countrycodes*

2 For a very good introduction to wind energy, try *http://www.windpower.dk/*.

3 Utility Wind-Modelling Planning Meeting (1996), National Renewable Energy Laboratory, Golden, Colorado, USA, 22-23 Feb 1996

4 Nielsen, L.H., and Morthorst, P.E. (ed.): *Fluktuerende vedvarende energi i el- og varmeforsyningen - det mellemlange sigt*. Risø-R-1055(DA), Forskningscenter Risø, Roskilde, April 1998, 154 pp. ISBN 87-550-2396-7 *(in Danish)*

5 Kalnay, E., S.J. Lord, and R.D. McPherson: *Maturity of Operational Numerical Weather Prediction: Medium Range*. Bulletin of the American Meteorological Society **79**(12), pp. 2753-2769 (1998)

6 *http://www.dwd.de/research/model/lm.html, 1999-12-03*

7 Haltiner & Williams: *Numerical Prediction and Dynamic Meteorology*. 2nd ed, Section 13-3

8 Lorenz, E.N.: *Deterministic Nonperiodic Flow*. J Atmospheric Sciences **20**, pp. 130-141 (1963)

9 Machenhauer, B. (Ed): *HIRLAM final report*, HIRLAM Technical Report 5, Danish Meteorological Institute, Copenhagen, Denmark (1988). For an introduction, see also *http://www.dmi.dk/eng/f+u/* .

10 *http://www.ecmwf.int/*

11 Mortensen, N.G., L. Landberg, I. Troen, and E.L. Petersen: *Wind Atlas Analysis and Application Program (WAsP) - User's Guide*. Risø-I-666(EN)(v.2), Risø National Laboratory, Roskilde, Denmark, 133 pp.

12 Troen, I., and E.L. Petersen: *European Wind Atlas*. Published for the EU Commission DGXII by Risø National Laboratory, Denmark (1998), ISBN 87-550-1482-8

13 Jackson, P.S. and J.C.R. Hunt: *Turbulent flow over a low hill*. Q. J. R. Meteorol. Soc. **101**, pp. 929-955 (1975)

14 Mason, P.J. and R.I. Sykes: *Flow over an isolated hill of moderate slope*. Q.J.R. Meteorol. Soc. **105**, pp. 383-395 (1979)

15 Taylor, P.A., J.L. Walmsley and J.R.Salmon: *A Simple Model of Neutrally Stratified Boundary-Layer Flow Over Real Terrain Incorporating Wavenumber-Dependent Scaling*. Boundary-Layer Meteorol. **26**, pp. 169-189 (1983)

16 Mason, P.J. and J.C. King: *Measurements and Predictions of Flow and Turbulence over an Isolated Hill of Moderate Slope*. Q.J.R. Meteorol. Soc. **111**, pp. 617-640 (1985)

17 Walmsley, J.L., P.A. Taylor, and T. Keith: *A Simple Model of Neutrally Stratified Boundary-Layer Flow Over Complex Terrain with Surface Roughness Modulations (MS3DJH/3R)*. Boundary-Layer Meteorol. **36**, pp. 157-186 (1986)

18 Troen, I.: *On Diagnostic Wind Field Models*. From College on Atmospheric Boundary Layer Physics, Trieste, Italy, 1990.

19 *E.g.* de Baas, A.: *The K-ε model in separating flows*. In: Proceedings of the Ninth Symposium on Turbulence and Diffusion, Roskilde. American Meteorological Society, pp. 359-362

20 Bowen, A.J., and N.G. Mortensen: *Exploring the Limits of WasP, the Wind Atlas Analysis and Application Program*. Proceedings of the European Union Wind Energy Conference held at Göteborg, Sweden, 20-24 May 1996, pp. 584-587, ISBN 0-9521452-9-4

21 Mortensen, N.G., and E.L. Petersen: *Influence of Topographical Input Data on the Accuracy of Wind Flow Modelling in Complex Terrain*. Proceedings of the European Wind Energy Conference held at Dublin, Ireland, October 1997, pp. 317-320, ISBN 0 9533922 0 1

22 Wood, N.: *The onset of separation in neutral, turbulent flow over hills*. Boundary-Layer Meteorology **76**, pp. 137-164 (1995)

23 Rao, K.S., J.C. Wyngard, and D.R. Coté: *The Structure of the Two-Dimensional Internal Boundary Layer over a Sudden Change of Surface Roughness*. J. Atmos. Sci. **26**, pp. 432-440 (1974)

24 Sempreviva, A.M., S.E. Larsen, N.G. Mortensen, and I. Troen: *Response of Neutral Boundary Layers to Changes of Roughness*. Boundary-Layer Meteorol. **50**, pp. 205-225

25 Sempreviva A.M. and A. Lavagnini: *Approccio modellistico alla climatologia del vento*. Convegno Societa' Italiana di Fisica, Sif, Padova 2-7 Ottobre 1986. IFA-C.N.R. Internal Report 86/39. *(in Italian)*

26 Krieg, R.: *Verifiering av beräknad vindenergiproduktion*. SMHI Rapport 1999 No. 28 *(in Swedish)*

27 Mortensen, N.G., P. Nielsen, L. Landberg, O. Rathmann and M.N. Nielsen: *A Detailed and Verified Wind Resource Atlas for Denmark*. Proceedings of the European Wind Energy Conference, Nice, France, 1-5 March 1999, p. 1161-1164, ISBN 1 902916 00 X

28 Sanderhoff, P.: *PARK - User's Guide, A PC-program for calculation of wind turbine park performance*. Risø-I-668(EN), 1993, Risø National Laboratory, Roskilde, Denmark, 8pp.

29 Jensen, N.O.: *A Note on Wind Generator Interaction*. Risø-M-2411, 1984, Risø National Laboratory, Roskilde, Denmark, 16pp.

30 Waldl, H.-P.: *Modellierung der Leistungsabgabe von Windparks und Optimierung der Aufstellungs-geometrie*. PhD-Thesis, Carl von Ossietzky Universität Oldenburg, 1997

31 Bossanyi, E.A.: *Short-Term Wind Prediction Using Kalman Filters*. Wind Engineering **9**(1), pp. 1-8 (1985)

32 *Informationen aus dem Forschungsschwerpunkt Energieversorgung mit dezentralen Kleinkraftwerken in leistungsbegrenzten Versorgungsnetzen.* Fachhochschule Wilhelmshaven, Fachbereich Elektrotechnik, Oktober 1999

33 Vihriälä, H., P. Ridanpää, R. Perälä, and L. Söderlund: *Control of a variable speed wind turbine with feedforward of aerodynamic torque*. Proceedings of the European Wind Energy Conference, Nice, France, 1-5 March 1999, pp. 881-884**,** ISBN 1 902916 00 X

34 Dambrosio, L, and D. Fortunato: *One-step-ahead control of a wind-driven, synchronous generator system*. Energy **24**, pp. 9-20 (1999)

35 Nogaret, E., G. Stavrakakis, J.C. Bonin, G. Kariniotakis, B. Papadias, G. Contaxis, M. Papadopoulos, N. Hatziargyriou, S. Papathanassiou, J. Garopoulos, E. Karagounis, J. Halliday, G. Dutton, J. Pedas-Lopes, A. Androutsos, and P. Pligoropoulos: *Development and Implementation of an Advanced Control System for Medium Size Wind-Diesel Systems*. Proceedings of the EWEC '94 in Thessaloniki, 10.-14. Okt, pp. 599-604

36 Tantareanu, C.: *Wind Prediction in Short Term: A first step for a better wind turbine control*. Nordvestjysk Folkecenter for Vedvarende Energi, October 1992, ISBN 87-7778-005-1

37 Dutton, A.G., G. Kariniotakis, J.A. Halliday, and E. Nogaret: *Load and Wind Power Forecasting Methods for the Optimal Management of Isolated Power Systems with High Wind Penetration*. Wind Engineering **23**(2), pp. 69-87 (1999)

38 Kariniotakis, G., E. Nogaret, and G. Stavrakis: *Advanced Short-Term Forecasting of Wind Power Production*. Proceedings of the European Wind Energy Conference held in Dublin, Ireland, October 1997, pp. 751-754, ISBN 0 9533922 0 1

39 Kariniotakis, G.N., E. Nogaret, A.G. Dutton, J.A. Halliday, and A. Androutsos: *Evaluation of Advanced Wind Power and Load Forecasting Meghods for the Optimal Management of Isolated Power Systems*. Proceedings of the European Wind Energy Conference, Nice, France, 1-5 March 1999, pp. 1082-1085, ISBN 1 902916 00 X

40 Beyer, H.G., T. Degner, J. Hausmann, M. Hoffmann, and P. Ruján: *Short Term Prediction of Wind Speed and Power Output of a Wind Turbine with Neural Networks.* Proceedings of the EWEC '94 in Thessaloniki, 10.-14. Okt, pp. 349-352

41 Tande, J.O., and L. Landberg: *A 10 sec. Forecast of Wind Turbine Output with Neural Networks*. Proceedings of the 1993 ECWEC in Travemünde, 8.-12. March 1993, pp. 774-777, ISBN 0-9521452-0-0

42 Alexiadis, M.C., P.S. Dokopoulos, H.S. Sahsamanoglou, and I.M. Manousaridis: *Short-Term Forecasting of Wind Speed and Related Electrical Power*. Solar Energy **63**, pp. 61-68 (1998)

43 Bechrakis, D.A. and P.D. Sparis: *Wind Speed Prediction Using Artificial Neural Networks*. Wind Engineering **22**(6), pp. 287-295 (1998)

44 Landberg, L.: *Short-term Prediction of Local Wind Conditions*. PhD-Thesis, Risø-R-702(EN), Risø National Laboratory, Roskilde, Denmark 1994, ISBN 87-550-1916-1

45 Landberg, L., and S.J. Watson: *Short-term Prediction of Local Wind Conditions*. Boundary-Layer Meteorology **70**, p. 171 (1994)

46 Joensen, A., G. Giebel, L. Landberg, H. Madsen, and H. Aa. Nielsen: *Model Output Statistics Applied to Wind Power Prediction*. Proceedings of the European Wind Energy Conference, Nice, France, 1-5 March 1999, pp. 1177-1180, ISBN 1 902916 00 X

47 Madsen, H. (Ed.): *Models and Methods for Predicting Wind Power.* ELSAM, Skærbæk, Denmark, ISBN 87-87090-25-2 (1995)

48 Nielsen, T.S., and H. Madsen: *Statistical Methods for Predicting Wind Power*. Proceedings of the European Wind Energy Conference held in Dublin, Ireland, October 1997, pp. 755-758, ISBN 0 9533922 0 1

49 Nielsen, T.S., H. Madsen, and J. Tøfting: *Experiences with Statistical Methods for Wind Power Prediction*. Proceedings of the European Wind Energy Conference, Nice, France, 1-5 March 1999, pp. 1066-1069, ISBN 1 902916 00 X

50 Mönnich, K.: *Vorhersage der Leistungsabgabe netzeinspeisender Windkraftanlagen zur Unterstützung der Kraftwerkseinsatzplanung*. PhD-thesis, Carl von Ossietzky Universität Oldenburg, 2000.

51 Beyer, H.G., D. Heinemann, H. Mellinghoff, K. Mönnich, and H.-P. Waldl: *Forecast of Regional Power Output of Wind Turbines*. Proceedings of the European Wind Energy Conference, Nice, France, 1-5 March 1999, pp. 1070-1073, ISBN 1 902916 00 X

52 Bailey, B., M. C. Brower, and J. Zack: *Short-Term Wind Forecasting*. Proceedings of the European Wind Energy Conference, Nice, France, 1-5 March 1999, pp. 1062-1065, ISBN 1 902916 00 X; see also *http://www.truewind.com/*.

53 Zack, J.W., M.C. Brower, B.H. Bailey: *Validating of the Forewind Model in Wind Forecasting Applications*. Talk on the EUWEC Special Topic Conference Wind Power for the 21st Century, Kassel, Germany, 25-27 Sept 2000

54 Törnevik, H.: private communication, 21.12.1999

55 Jacobs, A.: *KALCORR: a Kalman-correction model for real-time road surface temperature forecasting*. KNMI technical report TR-198, DeBilt, 1997, ISBN 90-369-2119-8

56 Martin, F., R. Zubiaur, P. Moreno, S. Rodriguez, M. Cabre, M. Casanova, A. Hormigo, and M. Alonso: *Operational Tool for Short Term Prediction Model of Energy Production in Wind Power Plants at Tarifa (Spain)*. Proceedings of the 1993 ECWEC in Travemünde, 8.-12. March 1993, pp. 802-803, ISBN 0-9521452-0-0

57 Perez, I.M.: private communication, 8.3.2000

58 Papke, U., A. Petersen, and V. Köhne: *Evaluation and Short-Time-Forecast of WEC-Power within the power grid of SCHLESWAG AG*. Proceedings of the 1993 ECWEC in Travemünde, 8.-12. March 1993, pp. 770-773, ISBN 0-9521452-0-0

59 Papke, U., and V. Köhne: *Pelwin -- ein Windleistungsprognosesystem zur Unterstützung des EVU-Lastverteilers*. In 2. Deutsche Windenergie-Konferenz 1994, Tagungsband Teil 1 (DEWEK '94), Wilhelmshaven, Deutschland, 1994. Deutsches Windenergie Institut GmbH (DEWI).

60 Landberg, L.: *A mathematical look at a physical power prediction model*. Wind Energy **1**, pp. 23-28 (1998)

61 Joensen, A.: *Short-Term Wind Power Prediction*. PhD-thesis, Institute of Mathematical Modelling, Danish Technical University, IMM-PHD-2000-74. Lyngby, Denmark (2000).

62 Glahn, H., and D. Lowry: *The Use of Model Output Statistics (MOS) in Operational Weather Forecasting*. Journal of Applied Meteorology **11**, pp. 1203-1211 (1972)

63 Landberg, L.: *Short-term prediction of local wind conditions*. Wind Engineering into the 21st Century, Proceedings of the Tenth International Conference on Wind Engineering, Copenhagen/Denmark, 21-24 June 1999, Larsen, Larose & Livesey (eds), Balkema, Rotterdam, ISBN 90 5809 059 0, pp. 1997-2003

64 Kalman, R.E.: *A New Approach to Linear Filtering and Prediction Problems*. Transaction of the ASME-Journal of Basic Engineering, 35-45 (March1960); see also e.g. Ljung, L. and T. Söderström, 1983, *Theory and Practice of Recursive Identification*, MIT Press, Cambridge/Mass.

65 Nelder, J.A., and Mead, R. 1965, Computer Journal, vol. 7, pp. 308–313.

66 Press, W.H., *et al*: *NUMERICAL RECIPES IN C: THE ART OF SCIENTIFIC COMPUTING.* Cambridge University Press 1992, (ISBN 0-521-43108-5); the online edition is available at *http://cfatab.harvard.edu/nr/ bookcpdf.html*.

67 Heitkötter, J. and D. Beasley, eds.: *The Hitch-Hiker's Guide to Evolutionary Computation: A list of Frequently Asked Questions (FAQ).* USENET: comp.ai.genetic. Available via anonymous FTP from *rtfm.mit.edu/pub/usenet/news.answers/ai-faq/genetic/* (2000) About 110 pages.

68 *http://cfatab.harvard.edu/nr/bookcpdf/c10-4.pdf*

69 Bossanyi, E.A.: *Use of a grid simulation model for longer-term analysis of wind energy integration.* Wind Eng. **7**, pp. 223-246. (1983)

70 Bossanyi, E.A., and J.A. Halliday: *Recent developments and results of the Reading/RAL grid simulation model*. Proc. 5th BWEA Conf., Reading (Cambridge University Press, Cambridge, 1983), pp. 62-74.

71 Dutton, A.G., and J.A. Halliday: *A Comparison of two logistic models for load and wind power scheduling in an island power station.* Proceedings of the EWEC '94 in Thessaloniki, 10.-14. Okt 1994, pp. 1074-1079.

72 Drezga, I., and S. Rahman: *Input Variable Selection for ANN-based Short-Term Load Forecasting*, IEEE Trans on Pow Syst. **13**(4), pp. 1238-1244 (Nov 1998)

73 Daneshdoost, M., M. Lotfalian, G. Bumroongit, and J.P. Ngoy: *Neural Network with fuzzy set-based classification for short-term load forecasting*. IEEE Trans on Pow Syst. **13(**4), pp. 1386-1391 (Nov 1998)

74 Khotanzad, A., R. Afkhami-Rohani, and D. Maratukulam: *ANNSTLF - Artificial Neural Network Short-Term Load Forecaster - Generation Three*. IEEE Trans on Pow Syst. **13(**4), pp. 1413-1422 (Nov 1998)

75 *http://www.mapp.org/publications/documents/EIA411/97%20IEA_L%26C/97411.pdf*

76 Landberg, L., S.J. Watson, J. Halliday, J.U. Jørgensen and A. Hilden: *Short-term prediction of local wind conditions*. Report to the Commission of the European Communities, JOULE programme, JOUR-0091-C(MB), March 1994

77 Landberg, L.: *The Availability and Variability of the European Wind Resource*. Int J Solar Energy **18**, pp. 313-320 (1997)

78 Taken from REISI: *http://www.iset.uni-kassel.de:888/reisi_dw.html*.

79 *IEA Electricity Information 96*. OECD Publications, Paris 1996, ISBN 92-64-15585-6

80 Nuclear Engineering Int, vol **43**, no 532, pp. 36-39, Nov 1998

81 Kalney, E. *et al*: *The NCEP/NCAR 40-year reanalysis project*. Bulletin of the American Meteorological Society **77**, (1996), pp. 437–471; See also *http://wesley.wwb.noaa.gov/reanalysis.html*

82 Nielsen, T.S., A. Joensen, H. Madsen, L. Landberg and G. Giebel: *A New Reference for Predicting Wind Power*. Wind Energy **1**, pp. 29-34 (1998)

83 Joensen, A.K.: *Models and Methods for Predicting Wind Power*. IMM-EKS-1997-17, Institute of Math. Modelling, The Technical University of Denmark, Lyngby, 1997. *(in Danish)*

84 Ensslin, C., B. Ernst, M. Hoppe-Kilpper, W. Kleinkauf, and K. Rohrig: *Online Monitoring of 1700 MW Wind Capacity in a Utility Supply Area.* Proceedings of the European Wind Energy Conference, Nice, France, 1-5 March 1999, pp. 444-447, ISBN 1 902916 00 X

85 Milligan, M.R., A.H. Miller, and F. Chapman: *Estimating the Economic Value of Wind Forecasting to Utilities*. Presented at Windpower '95, Washington, D.C., March 27-30, 1995. NREL/TP-441-7803. Get it from *OSTI.gov*.

86 Hutting, H.K., and J.W. Cleijne: *The Price of Large Scale Offshore Wind Energy in a Free Electricity Market*. Proceedings of the European Wind Energy Conference, Nice, France, 1-5 March 1999, pp. 399-401, ISBN 1 902916 00 X

87 Nielsen, L.H., P.E. Morthorst, K. Skytte, P.H. Jensen, P. Jørgensen, P.B. Eriksen, A.G. Sørensen, F. Nissen, B. Godske, H. Ravn, C. Søndergreen, K. Stærkind, and J. Havsager: *Wind Power and a Liberalised North European Electricity Exchange*. Proceedings of the European Wind Energy Conference, Nice, France, 1-5 March 1999, pp. 379-382, ISBN 1 902916 00 X

88 Sørensen, B. and P. Meibom: *Can Wind Power be Sold in a Deregulated Electricity Market?* Proceedings of the European Wind Energy Conference, Nice, France, 1-5 March 1999, pp. 375-378, ISBN 1 902916 00 X

89 Schwartz, M.N, and B.H. Bailey: *Wind Forecasting Objectives for Utility Schedulers and Energy Traders*. Presented at Windpower '98, Bakersfield, CA, April 27-May 1, 1998. NREL/CP-500-24680. Available from *OSTI.gov* .

90 Watson, S.J., G. Giebel, and A.K. Joensen: *The Economic Value of Accurate Wind Power Forecasting to Utilities*. Proceedings of the European Wind Energy Conference, Nice, France, 1-5 March 1999, pp. 1009-1012, ISBN 1 902916 00 X

91 Giebel, G., L. Landberg, K. Mönnich, and H.-P. Waldl: *Relative Performance of different Numerical Weather Prediction Models for Short Term Prediction of Wind Energy*. Proceedings of the European Wind Energy Conference, Nice, France, 1-5 March 1999, pp. 1078-1081, ISBN 1 902916 00 X

92 Joensen, A., L. Landberg, and H. Madsen: *A new measure-correlate-predict approach for resource assessment*. Proceedings of the European Wind Energy Conference, Nice, France, 1-5 March 1999, pp. 1157-1160, ISBN 1 902916 00 X

93 Landberg, L., M.A. Hansen, K. Vesterager, and W. Bergstrøm: *Implementing Wind Forecasting at a Utility*. Risoe National Laboratory, Roskilde, Denmark, March 1997. Risø-R-929(EN), ISBN 87-5502229-4

94 Ernst, B.: *Short-Term Power Fluctuations of Wind Turbines from the Ancillary Services Viewpoint*. Diplomarbeit, Institut für Solare Energieversorgungstechnik e.V., Querschnitts-Projektbereich Windenergie (Mittlerweile: Forschungsbereich Information und Energiewirtschaft). (1999) Kassel, Germany.

95 *Wind Force 10*. EWEA, Forum for Energy and Development, Greenpeace International, 1999, ISBN 1 871532 248

96 Milborrow, D.: *Capacity credits - clarifying the issues*. Proc. of the BWEA conference on Wind Energy 1996, p. 215-219

97 Kaltschmitt, M., and A. Wiese: *Zur Definition der Kapazitätseffekte einer Stromerzeugung aus Windkraft und Solarstrahlung*. BWK **48**, Nr. 7/8, 1996, p. 67-71 *(in German)*

98 Selzer, H.: *Wind Energy. Potential of Wind Energy in The European Community. An Assessment Study.* SOLAR ENERGY R&D IN THE EC SERIES G: Wind Energy, Volume 2. D. Reidel Publishing Company, Hardbound, February 1986, 160 pp., ISBN 90-277-2205-6

99 Diesendorf, M., B. Martin, and J. Carlin: *The Economic Value of Wind Power in Electricity Grids*. Proceedings of the International Colloquium on Wind Energy Brighton, UK, 1981, p. 127-132, ISBN 0-90608559-4

100 Johanson, E.E., and M.K. Goldenblatt: *An Economic Model to Establish the Value of WECS to a Utility System*. Second International Symposium on Wind Energy Systems, October 3rd-6th, Amsterdam, The Netherlands, p. G2-9 - G2-26, ISBN 0 906085 04 7

101 Bernow, S., B. Biewald, J. Hall, and D. Singh: *Modelling Renewable Energy Resources: A Case Study of Wind*. July 1994. ORNL/41X-03370V

102 *Wind Power Penetration Study, The case of Ireland*. Commission of the European Communities Report EUR 14250 EN, Brussels/Luxembourg 1992, 172 pp.

103 Gibbons, T.G., J. Haslett, E. Kelledy, and M. O'Rathaille: *The Potential Contribution of Wind Power to the Irish Electricity Grid*. Statistics and Operations Research Laboratory, Trinity College Dublin, September 1979

104 *Wind Power Penetration Study, The case of UK CEGB system*. Commission of the European Communities Report EUR 14245 EN, Brussels/Luxembourg 1992, 33 pp.

105 Lipman, N.H., P.J. Musgrove, P.D. Dunn, E. Bossanyi, G.E. Whittle, and J.A. Halliday: *Wind Energy Systems Integration Studies by the Reading University + Rutherford and Appleton Laboratries' Group*. Proceedings of the International Colloquium on Wind Energy Brighton, UK, 1981, p. 91-96, ISBN 0-90608559-4

106 *Wind Power Penetration Study, The case of Portugal*. Commission of the European Communities Report EUR 14247 EN, Brussels/Luxembourg 1992, 104 pp.

107 *Wind Power Penetration Study, The case of Denmark*. Commission of the European Communities Report EUR 14248 EN, Brussels/Luxembourg 1992, 124 pp.

108 *Vindkraft i Elsystemet*. Energiministeriets og Elværkernes Vindkraftprogramm, EEV 83-02, September 1983 *(in Danish)*

109 *Wind Power Penetration Study, The case of Greece*. Commission of the European Communities Report EUR 14252 EN, Brussels/Luxembourg 1992, 71 pp.

110 *Wind Power Penetration Study, The case of The Netherlands*. Commission of the European Communities Report EUR 14246 EN, Brussels/Luxembourg 1992, 46 pp.

111 Wijk, A.J.M. van, J.P. Coelingh, and W.C. Turkenburg: *Modelling Wind Power Production in The Netherlands*. Wind Eng. **14** no. 2 (1990), p. 122-140 (and references cited therein)

112 Halberg, N.: *An Assessment of the Large Scale Integration of Wind Power in The Netherlands*. In: Wind energy (Solar energy R&D in the European community. Series G; v. 1), D. Reidel Publishing Company, Dordrecht, Holland, 1982, p. 182-187, ISBN 90-227-1603-X; the quotation in the footnote in the introduction is cited on p. 160 as: Beitrag und Struktur regenerativer Energiequellen in der Bundesrepublik Deutschland 1980 und 2000, Schaubild 1.10a vom März 1982 des Foliendienstes Energiewirtschaft e.V., Bonn 1982.

113 *Wind Power Penetration Study, The case of the Federal Republic of Germany*. Commission of the European Communities Report EUR 14249 EN, Brussels/Luxembourg 1992, 68 pp. plus pictures.

114 Jarass, L.: *Strom aus Wind: Integration einer regenerativen Energiequelle*. Berlin, Heidelberg, New York: Springer, 1981. ISBN 3-540-10436 *(in German)*

115 *Wind Power Penetration Study, The case of Spain*. Commission of the European Communities Report EUR 14251 EN, Brussels/Luxembourg 1992, 81 pp.

116 *Wind Power Penetration Study, The case of Italy*. Commission of the European Communities Report EUR 14244 EN, Brussels/Luxembourg 1992, 73 pp.

117 Fischedick, M., M. Kaltschmitt: *Integration einer windtechnischen Stromerzeugung in den konventionellen Kraftwerkspark*. Proceedings of the 1993 ECWEC in Travemünde, 8.-12. March 1993, pp. 778-781, ISBN 0-9521452-0-0 *(in German)*

118 Kaltschmitt, M: *Möglichkeiten und Grenzen einer Stromerzeugung aus Windkraft und Solarstrahlung am Beispiel Baden-Württembergs*. Forschungsbericht des Instituts für Energiewirtschaft und Rationelle Energieanwendung, Stuttgart 1990 *(in German)*

119 Papadopoulos, M, P. Malatestas, S. Papathanassiou, and N. Bilios: *Impact of High Wind Penetration on the Power system of Large Islands*. Proceedings of the 1993 ECWEC in Travemünde, 8.-12. March 1993, pp. 782-786, ISBN 0-9521452-0-0

120 Saramourtsis, A.D., P.S. Dokopoulos, and I.M. Manousaridis: *Comparison of probablilistic and Mote Carlo techniques for evaluation of the performnce of wind-diesel energy systems*. Proceedings of the EWEC '94 in Thessaloniki, 10.-14. Okt, pp. 1092-1097

121 Hansen, J.C., and J.O.G. Tande: *Are Feasibility Studies Reliable at High Wind Energy Penetration Levels*. Proceedings of the 1993 ECWEC in Travemünde, 8.-12. March 1993, pp. 359-362, ISBN 0-9521452-0-0

122 Tande, J.O., and J.C. Hansen: *Wind Power Fluctuations Impact on Capacity Credit*. Proceedings of the European Union Wind Energy Conference held at Göteborg, Sweden, 20-24 May 1996, pp. 1089-1092, ISBN 0-9521452-9-4

123 Torre, M.C., P.Poggi, and A.Louche: *Integration Limit of Wind Turbines Generator in an Islander Electrical Grid - Case Study of Corsica.* Proceedings of the European Wind Energy Conference, Nice, France, 1-5 March 1999, pp. 923-926 , ISBN 1 902916 00 X

124 Solanky, B., A. Sharma, and T.K. Moulik: *Sustainable energy: 2012-policy and legislation.* IECEC-97. Proceedings of the Thirty-Second Intersociety Energy Conversion Engineering Conference (Cat. No.97CH36203), 27 July-1 Aug. 1997, Honolulu, HI, USA, Band 4 (1997) pp. 2345-2349, New York, NY, USA: IEEE, ISBN 0-7803-4515-0.

125 Lakkoju, V.N.M.R.: *Combined power generation with wind and ocean waves*, WREC 1996

126 Räsänen, M., J. Ruusunen, and R.P. Hämäläinen: *Customer level analysis of dynamic pricing experiments using consumption-pattern models*. Energy **20**(9), pp. 897-906 (1995)

127 Steinberger-Willms, R.: *Untersuchung der Fluktuationen der Leistungsabgabe von räumlich ausgedehnten Wind- und Solarenergie-Konvertersystemen in Hinblick auf deren Einbindung in elektrische Versorgungsnetze*. Dissertation an der Universität Oldenburg. Verlag Shaker, Aachen 1993, ISBN 3-86111-740-1, ISSN 0945-0726 *(in German)*

128 Milligan, M. and B. Parsons: *A Comparison and Case Study of Capacity Credit Algorithms for Wind Power Plants*. Wind Engineering **23**(3), 1999, p. 159-166. *Nearly identical to* CP-440-22591, available from *OSTI.gov*

129 Hirst, E., and B. Kirby: *Operating Reserves and Bulk Power Reliability*. Energy **23**(11), pp. 949-959 (1998)

130 Milligan, M.R.: *Variance Estimates of Wind Plant Capacity Credit*. AWEA Windpower '96, Denver, Colorado, USA, 23.-27. June 1996. NREL/TP-440-21311

131 Munksgaard, J., M.R. Pedersen, and J.R. Pedersen: *Samfundsmæssig værdi af vindkraft. Delrapport 1: Snæver samfundsøkonomisk vurdering.* AKF Forlaget, København, December 1995, ISBN 87-7509-429-0 *(in Danish)*

132 Nielsen, H.Aa., and H. Madsen: *Aggregated Power Consumption Models for the Eastern Part of Denmark*. Department of Mathematical Modelling, Technical University of Denmark, Lyngby 1997

133 Communication from the European Commission: *ENERGY FOR THE FUTURE: RENEWABLE SOURCES OF ENERGY, White Paper for a Community Strategy and Action Plan*, COM(97)599 final (26/11/1997), available from *http://www.cordis.lu/*

134 Bergmann, J.C.: *A Physical Interpretation of von Karman's Constant Based on Asymptotic Considerations - A New Value*, J. Atmos. Sci. **55**, pp. 3403-3405 (1998)