Technical University of Denmark

DTU

On the Predictability of Hub Height Winds

Draxl, Caroline; Giebel, Gregor; Hahmann, Andrea N.; Delle Monache, Luca

Publication date:
2012

Document Version
Publisher's PDF, also known as Version of record

Link back to DTU Orbit

DTU Library
Technical Information Center of Denmark

On the Predictability of Hub Height Winds

DTU Wind Energy
PhD-Report

Caroline Draxl
Risø-PhD-84(EN)
March 2012

DTU Wind Energy
Department of Wind Energy

On the Predictability of Hub Height Winds

Caroline Draxl

Author: Caroline Draxl
Title: On the Predictability of Hub Height Winds
Division: DTU Wind Energy

Abstract:

Wind energy is a major source of power in over 70 countries across the world, and the worldwide share of wind energy in electricity consumption is growing. The introduction of significant amounts of wind energy into power systems makes accurate wind forecasting a crucial element of modern electrical grids. These systems require forecasts with temporal scales of tens of minutes to a few days in advance at wind farm locations. Traditionally these forecasts predict the wind at turbine hub heights; this information is then converted by transmission system operators and energy companies into predictions of power output at wind farms. Since the power available in the wind is proportional to the wind speed cubed, even small wind forecast errors result in large power prediction errors. Accurate wind forecasts are worth billions of dollars annually; forecast improvements will result in reduced costs to consumers due to better integration of wind power into the power grid and more efficient trading of wind power on energy markets.

This thesis is a scientific contribution to the advancement of wind energy forecasting with mesoscale numerical weather prediction models. After an economic and theoretical overview of the importance of wind energy forecasts, this thesis continues with an analysis of wind speed predictions at hub height using the Weather Research and Forecasting (WRF) model. This analysis demonstrates the need for more detailed analyses of wind speeds and it is shown that wind energy forecasting cannot be reduced solely to forecasting winds at hub height. Calculating only the power output from hub height winds can result in erroneous estimates due to the vertical wind shear in the atmospheric boundary layer (PBL). Results show that the accuracy of modeled wind conditions and wind profiles in the PBL depends on the PBL scheme adopted and is different under varying atmospheric stability conditions, among other modeling factors. This has important implications for wind energy applications: shallow stable boundary layers can result in excessive wind shear, which is detrimental for wind energy applications. This is particularly relevant with offshore facilities, which represent a significant portion of new wind farms being constructed. Furthermore, a novel aspect to this study is the presentation of a verification methodology that takes into account wind at different heights where turbines operate.

The increasing number of wind farm deployments represents a novel and unique data source for improving mesoscale wind forecasts for wind energy applications. These new measurements include nacelle wind speeds and the turbines' angle of rotation into the wind (yaw angles). This thesis continues with an extensive description of this new data set and its challenges in data assimilation, focusing on data from the Horns Rev I wind farm. Since wind farm data are such a dense data set there is need to derive representative information from the measurements, i.e., thin the data. Different thinning strategies and their impact on improving wind forecasts for wind power predictions are investigated with the WRF Four-Dimensional Data Assimilation system. The median of the whole wind farm was found to be the most successful thinning strategy. Nacelle winds and yaw angles are a promising data set to improve wind predictions downstream of a wind farm as well as at the wind farm itself: Their impact lasted up to 5 hours and depends on time of the day, forecast lead time and weather situation.

The thesis is submitted to the Technical University of Denmark in partial fulfilment of the requirements for the PhD degree.

Dansk Resumé

Vind energi er en hovedkilde til energi i mere end 70 lande i verden og andelen af vind energi i forbruget af elektricitet er stigende. Indførelsen af signifikante mængder af vind energy i energinettet nødvendiggører vindforudsigelser på tidsskalaer mellem ti minutter og et par dage i forvejen i områder med vindmølleparker. Transmission System Operators og energifirmaer omregner normalt vinden i navhøjden til forudsigelse af den forventede production af elektricitet fra en vindmøllepark. Fordi kraften i vinden er proportional med vinden i tredje potens, fører endog små fejl i vindforudsigelsen til større fejl i energiprognosen. Præcise vindforudsigelser er derfor årligt potentielt millioner af dollars værd: gennem forbedret handel på energimarkederne og integrering af strøm i energinettet.

Denne afhandling er et videnskabeligt bidrag til forbedringen af vindenergiforudsigelser med mesoskala numeriske vejrprognose modeller. Efter en økonomisk og teoretisk oversigt af vigtigheden af vindenergiforudsigelser fortsætter afhandlingen med en analyse af vindhastighedsforudsigelser i navhøjden ved brug af vejrprognose modellen WRF (Weather Research and Forecasting). Denne analyse viser nødvendigheden af detaljerede analyser af vindhastigheder og indikerer at vindhastighedsforudsigelser ikke alene må begrænses til forudsigelser i navhøjden. Beregninger af energien alene baseret på vindhastigheder i navhøjden kan resultere i forkerte vurderinger som følge af vertikale vindvariationer i det atmosfæriske grænselag. Resultater viser at nøjagtigheden af de modellerede betingelser og vindprofiler i grænselaget afhænger af den anvendte parameterisering i grænselaget (mellem andre modelleringsaspekter) og at denne nøjagtighed varierer i forskellige atmosfæriske stabilitetetsbetingelser. Dette har betydningsfulde konsekvenser for vindenergi udnyttelse over vand, hvor mange af de nye vindmølleparker skal bygges; lave stabile grænselag kan resultere i uforholdsmæssigt store vertikale vind variationer, der er skadelige for vindenergi udnyttelse. Endvidere præsenteres en verifikationsmetode der tager hensyn til vinden i de forskellige højder hvor turbiner findes.

Med det støt stigende antal vindmølleparker, især offshore, er et nyt og enestående datasæt til forbedring af mesoskala vindforudsigelser gennem data assimilering til rådighed, nemlig: vindhastighedsmålinger fra navet i en vindmølle samt krøjning (yaw angle), der beskriver turbinens rotation med vinden. Efterfølgende fortsætter denne afhandling derfor med en omfattende beskrivelse af dette nye datasæt og dets udfordringer i data assimilering, fokuserende på Horns Rev I vindmøllepark. Fordi der er mange vindmøllemålinger på et begrænset område, data skal tyndes. Strategier for at tynde dem og værdien af dette nye datasæt i forbedringen af vindforudsigelser til vindenergi er undersøgt med WRF Four-Dimensional Data Assimilation system. Det vises at medianen af hele havmølleparken er den mest successrige tyndestrategi. Nacelle vinde og krøjning er et lovende datasæt til at forbedre vindprognoser nedstrøms fra vindmølleparken og på selve vindmølleparken: Deres effekt varer op til 5 timer og afhænger af tidspunkt om dagen, forecast lead time og vejret.

Acknowledgements

Coming together is a beginning. Keeping together is progress. Working together is success.

Henry Ford (1863 - 1947)

In that sense, I would like to thank those whose collaboration and assistance during the last 3 years made this PhD thesis possible.

I want to express my thanks to my supervisors Gregor Giebel, Andrea N. Hahmann and Luca Delle Monache for their strategic, organizational and scientific support, and their time.

Thanks to all my collegues from the Wind Energy Division at Risø DTU for the nice and interesting working environment, especially Mike Courtney (thanks for the motivating words!), Nicos Cutululis, Oliver Gehrke, Sven Erik-Gryning, Mark Kelly, Eleni Markou, Morten Nielsen, Søren Ott, Troels Friis Pedersen and Alfredo Diaz Peña. Anne Margrethe Larsen is thanked for her invaluable and immediate help in computing.

Thanks to Luca Delle Monache, Hans Huang, and Yubao Liu for hosting me at NCAR. I am also grateful to all who made my stay at NCAR welcoming and scientifically fruitful: Tom Auligne, Al Bourgeois, Barb Brown, Yaodeng Chen, Will Cheng, Gael Descombes, Yong-Run Guo, Jason Knievel, Branko Kosovic, Yuewei Liu, Linlin Pan, Ned Patton, Matt Pocernic, Syed Rafat Husain Rizvi, Gregory Roux, Francois Vandenberghe, Xiaoyan Zhang, and last but not least the DART group, expecially Jeff Anderson, Nancy Collins, So-Young Ha, Tim Hoar, David Powell, Glen Romine, and Chris Snyder.

A big thank you to Joseph Olson for discussions in boundary layer research and his help.

Many thanks to my office mates Ioanna Karagali, Joakim Refslund, Claire Vincent and Rozenn Wagner for their friendship, collaboration and the friendly atmoshpere! Claire is especially thanked for the proof-reading of my manuscripts, and Joakim for his help in translating the Danish abstract.

I am further grateful for inspiring and interesting discussions with Stan Benjamin, John Brown, Jochen Cleve, Nina Detlefsen, Josh Hacker, Rolf-Erik Keck, Lars Kruse, Julie Lundquist, Kaspar Mortensen, Jesper Nielsen Nissen, Pierre Pinson, Zhaoxia Pu and her research group, and Dusanka Zupanski.

Contents

CHAPTER 1

Introduction

Predicting the wind speed at the location of wind farms to estimate their power output is a major and important component in the integration of wind energy in power systems. In fact, the advancement of the wind energy sector and the increase of the wind energy share in the power mix within the last decades would not have been as successful without the use of wind forecasts.

Numerous efforts have been undertaken to predict the wind at hub height of large wind farms, ranging from statistical models for very short term forecasts of seconds and minutes in advance, to numerical weather prediction based models covering the medium and longterm forecasts of hours to days in advance.

Usually it is the wind at turbine hub height that is used by transmission system operators and energy companies to predict the power output. Predicting hub height winds with mesoscale weather prediction models is the focus of this thesis. The first part concentrates on the predictability of hub height winds with mesoscale models. The first question to be answered is: Is it enough to predict the wind at hub height to get accurate power estimates? Since wind turbines are situated in the planetary boundary layer (PBL), and the turbine rotors sometimes cover an area that is below and above boundary layer height if the boundary layer is shallow, calculating the power output from hub height winds only may result in erroneous estimates (Wagner, 2010). The first part of this thesis consists therefore of an examination of how well wind profiles can be predicted. A verification technique that takes the variation of wind with height into account is proposed.

Knowing about the deficiencies of a model is important - improved forecast models are needed, however, and could contribute to enormous savings in the wind energy sector. A major part of this thesis deals thus with the idea of improving meteorological forecasts. Much could be done in this respect - one way of improving the accuracy of numerical weather forecasts is through data assimilation. With a growing number of wind farm deployments, especially offshore, a lot of new observations are becoming available, since on every turbine the wind conditions are recorded for turbine control purposes. This information can be used in a data assimilation system to update the meteorological conditions at the site of interest,

a wind farm. Starting from updated initial conditions achieved through assimilating wind farm data will be shown to produce more accurate wind forecasts for a wind farm. The concept of assimilating wind farm data is novel terrain. Since current observations in the boundary layer are sparse and not always representative, wind farm data constitute a valuable additional data set. In this thesis, wind information collected from offshore wind turbine nacelles and angles of the turbine's rotation into the wind are assimilated into numerical weather prediction (NWP) models, with the aim of encouraging others to use wind farm data as well.

This thesis touches upon boundary layer research, specifically on winds in the boundary layer and on data assimilation in the boundary layer. As such it contributes to a better understanding of the wind conditions and the predictability of wind for wind farms. Since offshore wind farms are generally near the coast, the assimilation of wind farm data has the potential to improve wind and weather forecasts also inland.

Last but not least, this thesis is a contribution to the joint effort of everyone who promotes renewable energies.

This thesis is structured as follows: Chapter 2 introduces the importance of wind forecasts for wind energy purposes in detail. Their relevance in energy trading and grid integration is discussed. A section on wind energy forecasts places this thesis within the context of mesoscale NWP meteorological models with a special emphasis on wind forecasting.
Wind energy forecasting with mesoscale models is discussed in Chapter 3. I refer to the mesoscale Weather Research and Forecast (WRF) model which was used for the simulations in this thesis. After an initial evaluation of its hub height wind predictions, the focus is on the evaluation of winds and vertical wind shear from WRF model forecasts using seven boundary layer schemes. It is pointed out, that wind energy forecasting cannot be reduced to forecasting winds at hub height, nor at 10 m.
Since the initial conditions of a mesoscale NWP model determine, amongst other aspects, the forecast PBL structure and thus forecast winds, Chapter 4 deals with the idea of using data assimilation to improve wind energy forecasts for wind farms. The concept of data assimilation in general is described, followed by a section on data assimilation specifically for wind energy predictions.
With an increasing number of wind farm deployments, a novel and unique data set is becoming increasingly available: wind speeds measured on the nacelle of a wind turbine and the yaw angle, the turbine's angle of rotation into the wind. In Chapter 5 these wind farm observations are described. Since in the following chapter I will present results on the assimilation of these wind farm observations, chapter 5 also introduces the issues with these data in data assimilation.
Chapter 6 explores the benefits of the assimilation of wind farm measurements into the WRF Four-Dimensional Data Assimilation System (FDDA). First, aggregation strategies are presented on how to assimilate wind farm data; second, their assimilation impact is discussed. Finally, Chapter 7 concludes the thesis; an outlook and implications of the work are given in Chapter 8.

CHAPTER 2

Motivation

2.1   Importance of wind energy forecasting

More than 20% of the Danish domestic power consumption is supplied by wind power, which makes Denmark the leading country in terms of wind power penetration, followed by Spain, where the share of wind of electricity consumption is 12% (Krohn et al., 2009; Zervos and Kjaer, 2008). Introducing such significant amounts of wind energy into a power system entails both positive and negative economic impacts. In order to mitigate the negative impacts, wind power forecasts are important for mainly two main aspects: trading the power on energy markets and dealing with balancing power, which is essential for grid integration of wind power. Both aspects come down to the same priciple: demand and supply of power have to be balanced at all times.

Since the importance of wind forecasts can only be understood if the theoretical background of these two aspects is known, this chapter explains them from a perspective of liberalised EU power markets, with a special emphasis on the Nordic Power Market, which comprises Denmark, Norway, Sweden and Finland. The situation in the Nordic Power Market is driven by the fact that the produced wind power is given priority and has to be fed into the grid by law. Also in countries with an optional integration of produced wind power, e.g. in the U.S.A., there are incentives to use wind forecasts to determine operating reserve requirements (Botterud et al., 2009) or to trade the power on energy markets (Botterud et al., 2010; Marquis et al., 2011).

Placing this thesis within the perspective of these basic principles of wind power integration will conclude the chapter.

2.2 Trading wind power on energy markets

This section highlights the different horizons of wind energy forecasts and their importance for and relation to the power trade on energy markets.

2.2.1 Long horizon: 2-5 days

Wind power forecasts for the next 2 – 5 days, or even up to seven days, can give vital information about the expected market prices (risks of low prices, probability of high prices) and thereby when to make the outages for maintenance, operation tests, power tests, etc. This is the range where probabilistic forecasts, e.g., based on ensemble forecasts, can be very useful. This horizon is not covered in this thesis.

2.2.2 Medium horizon: 12-36 hours ahead

This is the horizon that is relevant for trading power on the spot market (day-ahead market). The spot market is a physical market, where prices and amounts are based on supply and demand. Power prices are determined with one price for each hour. Bidding closes at noon for deliveries from midnight and 24 hours ahead.

Accurate wind forecasts are important for this horizon for the following reasons:

- Estimation of the total wind power in the system and thus the power prices: The knowledge of the total wind power available for the day-ahead tells a lot about the market price for the day-ahead (Jónsson et al., 2010) and this information is vital for planning procedures of a power company. It is used to get a market overview. In order to determine the price for the day ahead, supply elements (like wind power, other power productions, power imports - the German power price is very sensitive to the wind power in Northern Germany, etc.) and demand elements (like consumption, which depends amongst other factors on weather and weekend/work day, and power exports) have to be balanced. Estimating the amount of wind power is linked to high uncertainty, whereas estimating the import/export volumes is linked to a moderate uncertainty for example. The rest of the elements like consumption or conventional production have a low uncertainty. This means that a good forecast of wind power allows market players to determine the price for the day ahead more accurately. Knowing the price for the day ahead, and since wind power in the energy mix reduces power prices in general (Krohn et al., 2009), electricity companies might use this information to deliberately shut down certain wind farms to keep the power prices at a certain level. Wind power reduces power prices, because the marginal price for wind is practically zero, since the prime mover - wind - is free.

- For trading on the day-ahead market: The better the forecast, the less power has to be traded on the intraday market, which will be shown to be more expensive in the next section.

2.2.3   Short horizon: 1-24 hours ahead

Most of the power of energy companies is traded on the day-ahead market. A dispatch operator of an energy company owning both wind and conventional power plants continuously makes sure that the total energy production is equal to the traded power. When wind power differs from the expected value, the operator must order the power plants to increase or decrease the production to keep the power balance. This is possible if regulation reserves are left. Otherwise, power will be sold/purchased on the intraday markets, which have a market gate closure of 1-2 hours before the delivery hour. The intraday market is a bilateral market, with different prices for the actors. The challenge on this market is to pick the best price at the best time. Any remaining imbalances will be traded on the regulating or balancing markets.

Both positive and negative imbalances are expensive for transmission system operators, energy providers and in the end for the whole society, and have to be sold/purchased on the balancing market. How expensive, depends on the general deficit/surplus of power on the market (Krohn et al., 2009): If the market tends towards a deficit of power, and if power production from wind power plants is lower than offered, other producers will have to adjust regulation up to maintain the power balance. The wind producers will be penalised and get a lower price for their electricity production than the spot market price. If wind power production is higher than the amount offered, wind power plants help to eliminate market deficits and therefore receive the spot price for the full production without paying a penalty. If the market tends towards an excess of power, and if power production from the wind power plant is higher than offered, other producers will have to adjust regulation down to maintain power balance. In this case, wind producers will be penalised and get a lower price for their electricity production than the spot market price. If the wind power production is lower than the bid, then wind power plants help to eliminate surplus on the market, and therefore receive the spot price for the full production without paying a penalty.

In this horizon, short-term wind forecasts are vital because frequently updated predictions help the dispatcher make the optimal choices on the intra-day market before it comes to be settled on the more expensive balancing markets. The better the forecasts, the smaller volumes will be exposed to these costly approaches.

Probabilistic forecasts are of no interest for the day ahead market, because the price determination is only affected by the bids and offers, and these are the average forecast volumes (without an indication of uncertainty).

2.2.4   Very short horizon: 10 Minutes to hours

This is a horizon relevant to power grid integration and is mostly covered by statistical models that include neural networks or are autoregressive moving average (ARMA) models. Since this horizon is not covered in this thesis, the interested reader is referred to Giebel et al. (2011) for an overview, or e.g., Gallego et al. (2011).

2.3 Impacts of wind power on the power system - coping with wind variability

Generally, the inherently variable nature of wind accounts for all the impacts and is the driver for the need of forecasts. Integrating wind into a power system requires changes in the scheduling and operation of other generators to deal with unpredicted deviations between supply and demand (Krohn et al., 2009).

Small temporal variations in wind energy production, like second to second or minute to minute, are rarely a problem for grid operators, since these variations will largely be cancelled out by other turbines in the grid. In general, the lowest time scale of interest to the Danish Transmission System Operator Energinet.dk is 5 minutes. Wind turbine energy production can, however, vary on an hourly basis, as can electricity demand from consumers. In both cases other generators on the grid have to provide power at short notice to balance supply and demand on the grid.

The cost of providing this balancing service depends on the type of generating equipment that is able to provide the balance supply and on the predictability of the variation in net electricity demand (demand variations minus wind power generation). The more predictable the net balancing needs, the easier it will be to schedule the use of balancing power plants and the easier it will be to use the least expensive units which are able to regulate power generation up or down at short notice.

It is not possible for wind producers to generate the amount of power they forecast on energy markets at all times, especially when bids have to be submitted between 12 – 36 hours in advance. Thus, other energy producers have to increase or reduce their power production to ensure that demand and supply are balanced. Also other actors on the spot market may require balancing power due to changes in demand, power plants shutting down, or other forced outages. Producers on the regulating market have to deliver their offers 1-2 hours before the hour of delivery, depending on the market rules. That is why only fast-response power producers (like gas fired power stations or hydro-electric power stations) will be able to respond quickly to deliver regulating power (Tuohy et al., 2009). The balancing costs, which can be significant, depend thus on the type and marginal costs of these reserve plants. Reducing these balancing costs is the main concern when integrating wind power in a power system.

The larger the geographical area to be balanced and the spread of wind power sites and aggregation in a system, the lower the effect of the wind variability on the grid (Landberg, 1997; Giebel, 2000; Focken et al., 2002). Larger balancing areas help decrease the impact of forecast errors of wind power, and thus reduce the amount of unforeseen imbalance (Krohn et al., 2009; Ernst et al., 2007; Giebel et al., 2007). Since offshore wind farms usually have an installed capacity that is larger than onshore wind farms, this effect is especially relevant for onshore wind farms. In the case of large offshore wind farms, like Horns Rev I and II with a high concentration of turbines in a limited geographical area, wind fluctuations can translate into large power fluctuations (Akhamatov et al., 2007; Sørensen et al., 2008). As wind farms get larger in the future and become more common, wind forecasts will be more and more important.

In summary, balancing costs could be decreased if wind forecasts were more accurate. Then only small deviations from the bidded wind power would be left to the rest of the power system.

2.4 Wind energy forecasts

Wind energy forecasts have been used and systems have been developed during the last decades, ranging from forecast times of minutes to hours or up to a few days, including statistical and meteorological models. Costa et al. (2008) give an extensive review on the history of wind power short-term prediction with a vast list of references. Within the ANEMOS.plus project, Giebel et al. (2011) published a comprehensive report on the State-Of-The-Art in Short-Term Prediction of Wind Power, stating that within the last 6 years, the field of short-term wind forecast has "practically exploded". The report covers time series models, meteorological models, short-term prediction, upscaling and spatio-temporal correlations, as well as ensemble forecasting to mention a few aspects. In these reports both power forecasting and wind forecasting are addressed.

The content of this thesis belongs to the field of mesoscale NWP meteorological models with a special emphasis on wind forecasting. As is inherent with these models, the forecast time scale is in the order of a few hours to a few days. More specifically, in the first part of the thesis I analyse forecasts ranging from 0 to 36 hours. The second part focuses on improving the short-term forecasts for wind farms, i.e., forecasts from 1 to 24 hours ahead. Improving the accuracy of these short-term forecasts has been shown in Costa et al. (2008), Giebel et al. (2011) and in Holttinen and Hirvonen (2005) to be of critical importance.

Through personal communication with transmission system operators and power companies, it seems that accurate *wind* forecasts are needed the most. The wind forecasts from wind at turbine hub height are then converted with in-house or commercial external tools to power forecasts for wind farms and balancing areas, which then serve as the basis for trading and scheduling. This thesis deals therefore with wind forecasts only. As further elaborated in Chapter 3, using predictions of winds at hub height only can result in erroneous power estimates (Wagner, 2010). In the future it will thus be necessary to use wind speeds at different heights to come up with a power estimate. An evaluation of mesoscale wind forecasts that takes winds at several heights into account is presented in section 3.3.

2.5 Summary and Conclusion

It has been argued that the use of wind forecasts is crucial to energy companies and transmission system operators in production planning and financial optimisation. Forecasts are needed for different time horizons and their primary benefit can mainly be grouped into the areas of grid integration and energy markets. The theories behind the economic benefits are very complex and depend on many aspects. Not two days are equal in assessing the value of improved forecasts. In this thesis I refer to *improved* wind forecasts in the sense of *more accurate* forecasts, i.e., a prediction of wind speed that is closer to the observed winds.

Whatever side you look at, it always comes down to financial benefits in the end. Short-term forecasts, i.e., forecasts up to a few hours ahead, seem to be among the most important when taking financial incentives as a benchmark to determine the importance of forecasts. Quantifying the monetary value of improved wind forecasts is complicated, however (Marquis et al., 2011). Studies conducted in the U.S.A. market seem to agree that the value is in the order of billions of dollars anually. Also in Europe, an improved quality of wind forecasts of only a few percent gives enormous benefits (Giebel et al., 2011).

If nothing else, improved grid integration of wind power and economic incentives through better forecasts enhance the value and acceptance of wind power, increase the share of renewable energy in the power mix, reduce the dependence of fossil fuels and thus help to mitigate the interference with climate and environment. This will ultimately be to the benefit of all of us.

Wind Energy Forecasting With Mesoscale Models

This chapter serves various purposes. It links the use of mesoscale models to wind energy forecasting and presents a first assessment of wind forecasts close to hub height modeled by the WRF model. This first assessment will demonstrate that verifying model winds at hub height only is not enough for wind energy purposes and motivates the principal part in this chapter: an analysis of seven boundary layer schemes of the WRF model, based on wind verification at different heights across the rotor area. A verification methodology of mesoscale model output for wind energy purposes is proposed. The evaluation of the seven boundary layer schemes and the new verification methodology are described in detail in Draxl et al. (2011b).

3.1   On the use of mesoscale models in wind energy forecasting

A mesoscale NWP model focuses per definition on the mesoscale, i.e., on a temporal scale of a few hours to a day and on a horizontal spatial scale of a few kilometers to several hundred kilometers (Pielke, 2002). The vertical scale extends from tens of meters to the depth of the troposphere. Pielke (2002) defines the mesoscale further as "those atmospheric systems that have a horizontal extent large enough for the hydrostatic approximation to the vertical pressure distribution to be valid, yet small enough for the geostrophic and gradient winds to be inappropriate as approximations to the actual wind circulation above the planetary boundary layer. This scale of interest, then, along with computer and cost limitations, defines the domain and grid sizes of mesoscale models." The hydrostatic approximation is good for synoptic and subsynoptic scales of motion and models using the hydrostatic approximation are usually applied to resolve resolutions as small as 10 km (AMS Glossary of Meteorology, http://amsglossary.allenpress.com/glossary/; last access Feb. 2012). These are mostly global and regional weather prediction models. For NWP models resolving small scale features in the order of tens of kilometers and small-scale mesoscale circulations such as

cumulus convection and sea-breeze circulations, the hydrostatic approximation is not made (non-hydrostatic models). These models require thus the full three-dimensional equations of motion.

Since the grid resolution in mesoscale models is too large to resolve explicitly the small-scale fluxes in the boundary layer, subgrid-scale turbulent fluxes and vertical mixing are handled by PBL parameterizations (Pielke, 2002; Stull, 1988a). These use the distribution of wind, temperature, and water vapor mixing ratio with height as well as the surface fluxes calculated from a land surface model and/or surface layer scheme to determine, amongst other things, the time tendencies of wind. The latter can then be used as predictions for wind energy purposes.

The Weather Research and Forecasting (WRF) model

The mesoscale model applied in this thesis for all the simulations is the Advanced Research WRF model (Skamarock et al., 2008), which is a community model maintained by the National Center for Atmospheric Research (NCAR) in the U.S.A. Its principal components comprise a dynamics solver, physics packages, and initialization routines. Multiple numerics/dynamics options can be handled by a namelist. It also includes a variational data assimilation (3DVAR and 4DVAR) package and nudging capabilities. The equations are fully compressible and Euler nonhydrostatic. The vertical coordinates are terrain-following, dry hydrostatic-pressure with vertical grid stretching permitted.

3.2   Assessing hub height wind forecasts in WRF

When using a NWP model, it is important to understand its deficiencies and physical limitations. For that purpose, I evaluated Risø's real-time WRF-based forecast system for the period from 1 March to 15 May 2009 for a first assessment (adapted from Draxl et al. (2009)).

3.2.1   WRF setup

Real-time test runs using the Advanced Research WRF Version 3.0.1.1 were in operation from March - May 2009 with the following specifications:

- Mother domain with 18 km resolution; two nested domains (two-way nests) with 6 and 2 km horizontal resolution (Figure 3.1);

- 37 vertical levels; the lowest levels were at approximate heights of 14, 55, 105 and 164 m AGL

- Initial and boundary conditions from US National Centers for Environmental Prediction (NCEP) Global Forecast System analyses at $1° \times 1°$ resolution; Sea surface temperature fields at high resolution from NCEP analysis

- Runs start at 6:00 and 18:00 local time from 00:00 and 12:00 UTC initial fields

Figure 3.1: WRF model domain configuration and surface elevation (left). Land-use classes on domain 3 (right); The WRF grid squares closest to those at Høvsøre and covering most of Denmark are classified as *Dryland cropland and pasture*. Also, location of observation masts at Høvsøre and the offshore wind farm Horns Rev.

- Model physics: Thompson graupel scheme, Kain-Fritsch cumulus parameterization. Lower atmosphere: Unified Noah landsurface model, Monin-Obukhov scheme, Yonsei University (YSU) PBL, 6th order numerical diffusion, Positive definite advection of moisture and scalars

- No FDDA or grid nudging

3.2.2 Verification procedure

The model was evaluated at a height of ~ 55 meters (second model level) for a period from 1 March to 15 May 2009. Winds at this height are compared to those measured with a Risø cup anemometer (Pedersen, 2004) at a height of 60 meters at Høvsøre (Figure 3.1) for the same time period. Since the measurement mast at Høvsøre is situated in the wake of wind turbines when the wind comes from northern directions, I excluded winds between 300 and $60°$ from the computations. In order to verify the forecast runs against the observations and to see a pattern of forecast skill, a poor man's ensemble was created by using different forecast runs for the same forecast time. Since the model was run every 12 hours for a period of 24 hours, it was thus possible to get up to three different forecasts for one specific time. These values were taken to compute an average bias and root mean square errors (RMSE) for every forecast hour (i.e. +1h, +2h, +3h, until +23h).
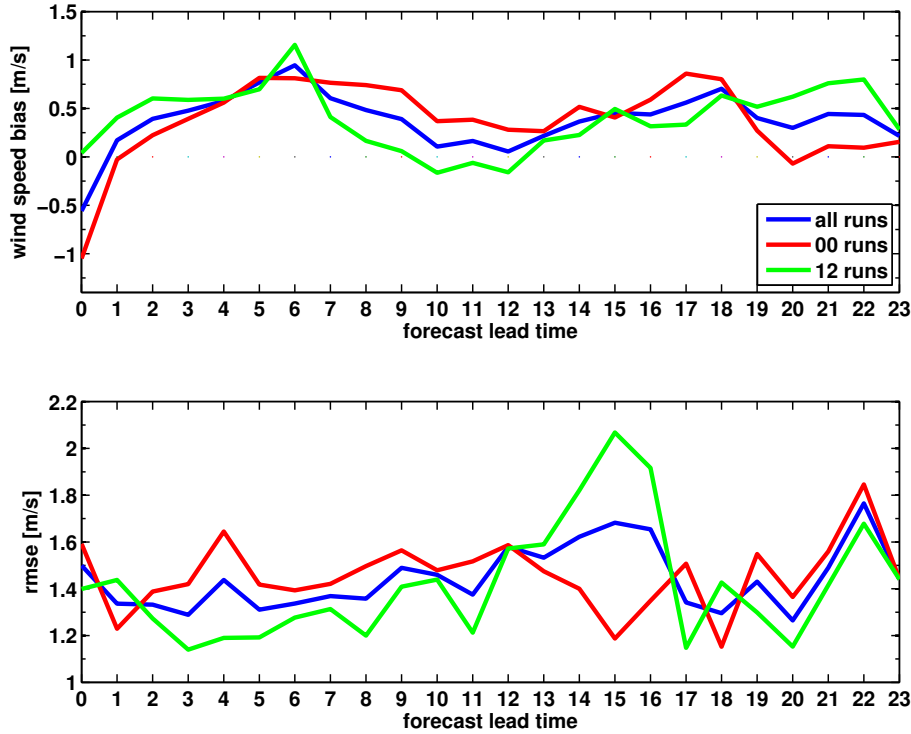
Figure 3.2: Wind speed biases (upper panel) and RMSE (lower panel) for Høvsøre for forecast lead times from 0 (analysis) to +23 hours for the WRF simulations initialized at 00 UTC and 12 UTC, respectively. The blue line depicts the bias for all the simulations.

3.2.3 Verification results

The mean root mean square errors and mean bias of wind speed for forecast lead times up to +23 hours are within $2 \, \mathrm{m\,s}^{-1}$ and approx. $1 \, \mathrm{m\,s}^{-1}$, respectively (Figure 3.2). Figure 3.2 (upper panel) shows a definite relation of bias to the time of the day. After some noise in the analysis, the forecasts are statistically more accurate for forecast times of +1 hour. This is especially true for the runs initialized at 00 UTC. The runs have their statistically worst time around a forecast time of +6 and +18 hours, showing more accurate behavior around leadtimes of +10 to +13 hours. This diurnal cycle can be explained by the model's failure of capturing transitions from stable to convective boundary layers: In the morning transition, that in Denmark in this season takes place usually around 6 UTC with sunrise, it is assumed that the nocturnal inversion gradually weakens due to warming and the convective boundary layer grows. In the evening transition, in Denmark around 18 UTC, the convective boundary layer dissolves and a surface inversion is formed. The +6 hour and +18 hour forecasts correspond to 6 and 18 UTC for the 00 UTC runs, and 18 and 6 UTC for the 12 UTC runs, respectively. That is the time of a transition that statistically is not perfectly captured and results thus in a higher bias. The +12 hour and +24 hour forecasts correspond to 12 and 00 UTC, respectively. No transition occurs in general at these times, so the bias gets less pronounced. There is a difference in the behavior of the runs initialized at 00 and 12 UTC. While the night runs start with a more negative bias for the analysis, the day runs start with

Figure 3.3: Wind speed biases (upper panel) and RMSE (lower panel) in $\mathrm{m\,s^{-1}}$ for Høvsøre for forecast lead times from 0 (analysis) to +23 hours for wind speed bins of 0–2.5 $\mathrm{m\,s^{-1}}$, 2.5–5 $\mathrm{m\,s^{-1}}$, 5–10 $\mathrm{m\,s^{-1}}$ and >10 $\mathrm{m\,s^{-1}}$ for the runs initialized at 00 and 12 UTC combined. The blue line depicts the bias and RMSE for all the simulations.
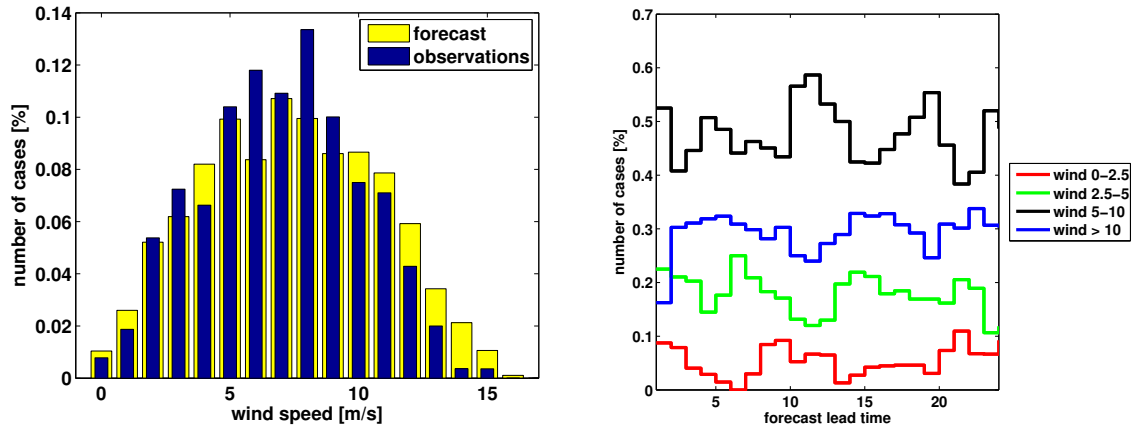
Figure 3.4: Left: Histogram of forecast (yellow) and observed wind speeds (blue). Right: Distribution of the different wind speed classes for each forecast hour.

a statistically more accurate forecast. I propose two reasons for the difference in the two cycles: (i) the runs initialized at night have fewer observations included in the analysis than

during the day, therefore the analysis values of the 12 UTC runs are more accurate; (ii) the large scale model's boundary layer used for the initialisation of the forecasts captures the nocturnal PBL worse. The root mean square error (Figure 3.2, lower panel) shows a quite homogeneous behavior for leadtimes up to +13 hours, followed by a striking peak of a higher error for the 12 UTC run. This behaviour would be smoothed out when using a verification period longer than nine weeks, since the statistics in that case are strongly influenced by a single poorly forecast event. For all runs it can be summarized, that WRF overpredicts the wind speeds slightly, having a positive bias most of the time (Figure 3.2 (upper panel) and Figure 3.4). The reason for overpredicting the wind speeds is probably because the model is prone to mix winds from above the boundary layer into stable atmospheres (Mass et al., 2002).

3.2.4 Dependence of error metrics on wind speed

I further calculated bias and RMSE values depending on the forecast wind speed. The forecast wind speeds were gouped into bins of 0–2.5 m s^{-1}, 2.5–5 m s^{-1}, 5–10 m s^{-1} and > 10 m s^{-1}. Figure 3.3 (upper panel) shows a clear dependence of bias on the forecast wind speed. The higher the forecast winds, the more positive the bias, e.g., if the model predicts wind speeds above 10 m s^{-1}, it is rather overpredicting, whereas when the model predicts low winds (0-5 m s^{-1}), it is under predicting in general. For winds > 2.5 m s^{-1}, the bias increases compared to the analysis for the next +4 to +6 forecast hours, rising afterwards. Since the amount of data is less in lower wind speed classes as depicted in Figure 3.4 (right), their bias curves are less smooth and single events have a larger impact on the statistics.

The root mean square error is less for low winds than for high winds (Figure 3.3, lower panel). One might argue that this behavior is an artifact of wind speeds being positive and because the possible span of predicting wind speeds wrongly gets higher at higher wind speeds. Statistics depend on their application, however. In wind energy, low winds < 2.5 m s^{-1} are not relevant. The higher the wind speeds, the bigger the impact of forecast errors, since the power in the wind depends on the wind speed cubed. This is especially important in the steep part of the power curve, i.e., depending on the turbine manufacturer between $\sim$3 and 15 m s^{-1}. Above 15 m s^{-1}, in the flat part of the power curve, small forecast errors are mostly inconsequential, except for predictions around cut off wind speeds at approx. 25 m s^{-1}. Thus, for wind energy purposes, it is often relevant to not flatten statistics but to explicitly look for single poor forecast events.

3.2.5 Summary

Although the data set used in this analysis is too short to draw overall conclusions of model behavior, it serves well to get a good overview of how the WRF model predicts wind speeds around hub height. This first assessment of the WRF model shows, that the wind forecasts at 60 m at Høvsøre show in general low root mean square errors and biases. However, overall statistics blur or flatten strong forecast differences that may occur in individual cases. Especially in wind speed classes that are crucial for no/full wind, like around the start up wind speed at approx. 3 m s^{-1} and the cut off wind speed at approx. 25 m s^{-1}, a predicted difference of only a few m s^{-1} can result in an important difference in power output. This

is most relevant around the cut off wind speed, since high wind power production rates are at stake. WRF tends to overpredict the wind speeds at 60 m, which could also be found in Draxl et al. (2011b). The higher the forecast winds, the more pronounced this effect.

It could further be shown, that the first forecast hours lack accuracy due to model inadequacies at the beginning of each run. The model starts from dry initial conditions only with moisture fields, i.e., all cloud-related fields are zero. In addition, initial fields are smooth because of the coarse resolution of the driving global model. Consequently, mesoscale features take some time to develop. The differences in behaviour between the 00 UTC and 12 UTC run, that I relate to a better analysis from the GFS model that is available at 12 UTC, shows the need to improve the initial conditions. This will will be further elaborated in chapter 4.

3.3   Evaluating winds and vertical wind shear from WRF model forecasts

Even though a first assessment is very helpful in understanding the model behavior, wind energy forecasting cannot be reduced to forecasting winds at hub height only or even at 10 m, as it is sometimes done. Wind speeds change within the interval of a rotor diameter of large wind turbines (typically between 30 and 130 m). This is especially important to consider in regions with shallow stable boundary layers, like the North Sea, where the boundary layer top can be around hub height (Giebel and Gryning, 2004). The lower part of the rotor will then be below and the upper part above the boundary layer top, leading to extreme wind shear across the rotor area. Considering the vertical variation of wind speeds with height (i.e., wind shear) in wind energy forecasting is crucial for several reasons:

- Wind shear is responsible for the development of turbulent motions associated with shear-induced instability (Kelley et al. (2004), Blumen et al. (2001)). These turbulent motions can lead to vibration or damaging loading events and thus to shorter life times of wind turbines.

- Wind shear impacts the extraction of energy from the wind over the area of the rotor disk (Wagner, 2010).

- The wake generated by a wind farm is affected by atmospheric stability conditions (Christiansen and Hasager, 2005a). Information about the wind shear, which is an indication for atmospheric stability, is important when assessing the overall power losses by a wind farm.

- When the output from mesoscale models is used for wind resource assessment, information about the wind profile across the rotor area will lead to more accurate resource estimates.

Draxl et al. (2011b) presents a more extensive comparison as done before of seven PBL parameterizations available in WRF V3.1 and focuses on the novel approach of analyzing not only wind speeds at hub height, but wind shear within that part of the boundary layer that covers the entire rotor of most modern and large wind turbines. Most of this paper will be shown in this section. Its purpose is twofold:

- By analyzing seven boundary-layer schemes of the WRF model, recommendations are provided about which PBL parameterization performs best for wind energy purposes under similar climatological and geographical conditions as the ones from this study.

- A verification methodology for mesoscale models used for wind energy applications is presented that takes winds at several heights into account.

3.3.1 Introduction

As mentioned in section 3.1, the choice of PBL scheme plays a significant role in the evolution of the low-level wind structure and therefore can heavily impact the quality of the forecast winds. Many aspects of a mesoscale model determine the quality of the forecast PBL structure. In the WRF model, surface momentum, heat and moisture fluxes are calculated by a surface layer scheme, that is coupled to a land surface model, which in turn provides input for the PBL scheme. Since each PBL scheme is expected to run with a particular surface layer scheme, an evaluation of the performance of the different PBL schemes alone is difficult (Shin and Hong, 2011). And, from a practical point of view, these differences are irrelevant due to the tight coupling between surface layer and PBL scheme. In this study, differences between the PBL schemes will be shown while evaluating properties of the surface layer and land surface models as well.

The verification site used in this study, Høvsøre, is situated near the coast. Although I chose the nearest model grid point over land to be compared with the measurements, it is likely there are influences of internal boundary layers, low level jets and sublayer structures not captured by the model resolution. In fact, the verification of these phenomena is an important research area in itself.

Others have used the WRF model to analyze wind predictions: Mesoscale model simulations using five of the PBL parameterizations were compared previously in Zhang and Zheng (2004) for surface wind and temperature in the Mesoscale Model Version 5 (MM5), and in Shin and Hong (2011) in WRF. Previous studies tackled the comparison with fewer schemes in WRF with a different focus: Li and Pu (2008) and Nolan et al. (2009) studied sensitivity during hurricane events, Jankov et al. (2005) discussed different WRF configurations in rainy conditions, Hu et al. (2010) focused on three PBL schemes; their wind profile analysis did not meet the required resolution for wind energy purposes, however. A vast comparison of different PBL parameterizations, using single column models, was carried out during the GEWEX Atmospheric Boundary Layer Study (Bosveld et al., 2008; Cuxart et al., 2006). Studies related to wind energy comprise a comparison of two boundary layer schemes with an emphasis on forecasting low-level jets (Storm et al., 2008), as well as an analysis of the wind shear over the United States Great Plains with four PBL parameterizations (Storm and Basu, 2010).

The novel approach in this paper consists of a more extensive comparison than done before, using seven PBL parameterizations available in WRF V3.1. I focus on analyzing wind shear within that part of the boundary layer that covers the entire rotor of most modern and large wind turbines, and I consider atmospheric stability.

3.3.2 Data

3.3.2.1 WRF Model setup

The WRF model setup (Advanced Research WRF, Version 3.1 (Wang et al., 2010)) consists of a main grid with horizontal grid spacing of 18 km and 2 nested domains (with 6 and 2 km spacing), the innermost domain covering most of Denmark (Figure 3.5). The model was initialized and forced at the boundaries by $1° \times 1°$ US National Center for Environmental Prediction (NCEP) Global Forecast System analyses at 6 hours intervals. Hence, the model simulations are actually hindcasts rather than forecasts. The sea surface temperature fields are also obtained from NCEP analyses at a horizontal resolution of $0.5° \times 0.5°$. Land use categories come from the United States Geological Service (USGS). I used 2-way nesting between domains and 37 vertical levels, with 8 levels within the lowest 500 m. The lowest levels important for wind energy applications were at approximately 14, 53, 105, and 164 m AGL during the studied period. The model physics options included: Thompson microphysics scheme, Kain-Fritsch cumulus parameterization, 6th order numerical diffusion, and positive definite advection of moisture and scalars. The rapid radiative transfer model (RRTM) and Dudhia schemes are used for longwave and shortwave radiation calculations, respectively. No data assimilation or grid nudging was used in the forecasts. These choices were based on experience from previous modeling systems Liu et al. (2008) and short sensitivity experiments. A similar model setup has been used in real-time WRF forecasts at DTU Wind Energy since May 2009.

Seven experiments with 30 hour forecasts were carried out for 1–30 October 2009. Each forecast was initialized at 12:00 UTC (UTC = local standard time - 1 hour). However, for the comparison only forecasts with lead times of 7 to 30 hours were used to avoid using forecasts within the model spin-up period. The seven experiments differ only in the PBL schemes, their associated land surface models and surface layer physics as recommended in Wang et al. (2010), and are described in Table 3.1. All other model aspects, such as convective and radiation parameterizations, remained the same.

The month of October 2009 was the period of choice due to its variable and representative weather conditions: The synoptic situation in Denmark was characterized by a low pressure system over Scandinavia during the first few days, interrupted by a ridge of high pressure on the 5th of October. From the 9th October, and for about 10 days onwards, anticyclonic conditions prevailed, leading to stable conditions at night and unstable conditions during the day. After that, an upper-
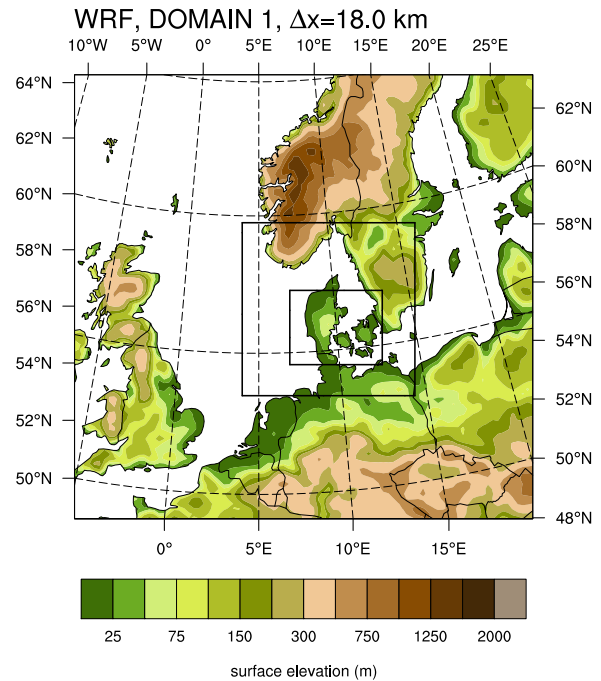

Figure 3.5: Domain configuration and terrain elevation of the WRF model setup. The black squares indicate the boundaries of two nested domains.

Table 3.1: Description of the seven experiments: PBL parameterizations, their closure type (Turbulence Kinetic Energy; TKE) associated land surface models and surface layer physics schemes, as recommended in Wang et al. (2010).

Experiment	PBL parameterization	Closure type	Land surface model	Surface layer scheme
ACM2	Asymmetric Convective Model version 2 Pleim (2007)	First Order Closure	Pleim-Xu	Pleim-Xu
MRF	Medium Range Forecast Model Troen and Mahrt (1986)	Non-local-K mixing	Unified Noah LSM	Monin-Obukhov
MYJ	Mellor-Yamada-Janjic Mellor and Yamada (1982)	TKE 1.5-order	Unified Noah LSM	Eta similarity
MYNN2	Mellor-Yamada Nakanishi and Niino Level 2.5 Janjic (2001)	TKE 1.5-order	Unified Noah LSM	MYNN
MYNN3	Mellor-Yamada Nakanishi and Niino Level 3 Nakanishi and Niino (2006)	TKE 2nd-order	Unified Noah LSM	MYNN
QNSE	Quasi-Normal Scale Elimination Sukoriansky et al. (2006)	TKE 1.5-order	Unified Noah LSM	QNSE
YSU	Yonsei University Scheme Hong et al. (2006)	Non-local-K mixing	Unified Noah LSM	Monin-Obukhov

level low over western Europe with neutral and slightly unstable conditions determined the weather conditions over Scandinavia again, followed by a high pressure system and characterized by stable conditions during night times by the end of the month. Stable atmospheric conditions in the lower boundary layer occur often in Denmark (Peña, 2009). In response to these large-scale conditions, the low-level flow at the verification site is mainly northwesterly (from the sea) during the first half of October and easterly (from land) during the second half.

3.3.2.2 Observations

The WRF model simulations were verified against measurements from a meteorological mast and a light tower at Høvsøre, situated on the northwest coast of Denmark, 1.7 km inland (Figure 3.6). DTU Wind Energy manages the National Test Station for Large Wind Turbines there and up to 5 wind turbines are in testing at any particular time. The terrain around the site is flat and homogeneous and the prevailing wind directions are west and northwest. In this study measurements of 10 min averaged wind speeds at heights of 10, 40, 60, 80, 100 and 116 m from the meteorological mast and at 160 m from the light tower are used; further, the temperature measurements at 2 and 100 m as well as kinematic heat fluxes measured by a sonic anemometer at 10 m. The friction velocity u_* and the Obukhov length L were computed from the sonic fluxes. When winds are from the north ($330°$–$30°$), the mast is located in the wake of a row wind turbines (Figure 3.6). Data within this range of directions were not used for the results presented here. The percentage of wake-free data in each atmospheric stability class (defined within intervals of L as proposed by (Gryning et al., 2007) relative to all the data per stability class varies from $72 - 93$ % (Table 3.3) during the time of the numerical experiments.

3.3.2.3 Verification methods

In this paper a verification method for mesoscale models for wind energy applications is proposed. Since the hub height of most large wind turbines is between 60 and 100 m, a

Figure 3.6: (left) Map of Høvsøre, Denmark. The large dots show the location of the wind turbines, squares show the location of the meteorological mast and light tower. The lines indicate the wake zone of the wind turbines (330°–30°). The inset shows Denmark; the arrow points at Høvsøre. (right) Location of Høvsøre (red plus sign) relative to the WRF model grid (2 km grid spacing). The green cross shows the location of the model grid point used for verification. Ocean grid squares are blue; yellow grid boxes are dryland cropland and pasture; the brown box is Cropland/Grassland Mosaic.

mesoscale model must be verified at these heights at least. Furthermore, information about the wind profile is crucial for wind energy forecasting because rotor diameters of large wind turbines are between 30 and 130 m. Within this interval, the turbines might experience a wind shear across the rotor area in excess of 5 m s^{-1}. This shear is often observed in regions with shallow stable boundary layers, like in the North Sea, where the boundary layer top can be around hub height (Giebel and Gryning, 2004). The lower part of the rotor will then be below and the upper part above the boundary layer top.

In the comparison I concentrate thus on verifying the model forecasts at different model levels (section 3.3.3.1) and the wind shear (section 3.3.3.2) and its dependence on atmospheric stability conditions (section 3.3.3.3 and Table 3.3). Error metrics such as the root-mean-square-error and a newly defined "wind profile error" (section 3.3.3.4) lead to the conclusions.

Since the WRF mass grid point closest to Høvsøre is classified as water (Figure 3.6b), the next grid point eastward is chosen for comparison against observations. This grid point displays similar distance to the coast as in reality and is classified as "Dryland cropland and pasture", with a surface roughness length of 0.15 m during 1–14 October and 0.05 m after. The observational estimates of surface roughness length at the site are in the range of 0.01–0.02 m all year round (Gryning et al., 2007; Peña, 2009).

The power exponent parameter α defined from the power law is used:

$$\frac{u_2}{u_1} = \left(\frac{z_2}{z_1}\right)^{\alpha} \tag{3.1}$$

where u_1 and u_2 are the wind speeds at levels z_1 and z_2, respectively, to evaluate the wind shear in the observations and model simulations.

An estimate of the value of α in the surface layer can be calculated from Monin-Obukhov similarity theory (MOST, e.g., Stull (1988b)). For a given set of heights, the value of α depends on atmospheric stability and roughness length. In the calculations, z_1 and z_2 represent the measurement heights of 10 and 60 m. For the model data, the wind speed at the first and second model levels are used and their respective heights at approximately 14 and 53 m. Replacing u_1 and u_2 in Eq.3.1 with the logarithmic wind law and for a surface roughness of $z_0 = 0.02$ m, which is close to the observed value at Høvsøre (Peña, 2009), $\alpha \simeq 0.14$. Based on the roughness length description of the WRF setup and under neutral conditions, the values of alpha are slightly larger: $\alpha \simeq 0.16$ and $\alpha \simeq 0.19$ for $z_0 = 0.05$ m (winter value) and $z_0 = 0.15$ m (summer value), respectively.

Gryning et al. (2007) showed, using data from Høvsøre, that on a semi-log graph the wind profile in statically neutral situations ($|L| > 500m$) appears as a straight line, whereas wind profiles in stable ($L < 500$ m) and unstable ($L < -500$ m) boundary layers curve downward/negatively and upward/positively with height. Model-derived wind profiles are thus compared under distinct classes of L to evaluate the ability of the various PBL schemes to represent such behavior. The wind profile error, WPE, is calculated by averaging the root-mean-square errors for the observed wind speeds at 60, 100 and 160 m and the second, third and fourth model level (at 53, 105, and 164 m).

3.3.3 Verification of low-level winds

The simulated wind speed was verified against observations at model levels one to four; observed temperatures against those at 2 m and on the second model level. No vertical interpolation of the model values to match the measurement heights has been done since the height of the model levels is often quite close to that of the observations (only at the second level the average difference in height is 7 m). In addition, I used the observed values at Høvsøre of the shear exponent, α, the friction velocity, u_*, the kinematic heat flux at 10 m, $\overline{w'T'}$, and the Obukhov length, L. From the measurements, $L = -(u_*^3 T_0)/(\kappa g \overline{w'T'})$. T_0 is a fixed temperature of 20°C, and κ and g are the von Kármán constant ($\kappa = 0.4$) and the acceleration due to gravity, respectively. For the model, the same relationship is used but with the model-derived values that are instantaneous and output hourly.

The overall error statistics between the simulated values from the seven sets of forecasts and the observations are presented in Table 3.2. Biases in wind speed are mostly positive (i.e., the simulated wind speeds are higher than those observed), except for the QNSE- (at 10 and 60 m), the MYJ- (at 60 m) and YSU-based (at 160 m) forecasts that have small negative biases. The absolute values of the biases are larger at the first model level than at 160 m. Smaller errors at 10 m are sometimes but not always correlated to smaller errors at 60 and 160 m. Temperature biases are within ± 1°C, except for the YSU-based forecasts at 2 m. Thus, the forecasts tend to be slightly warmer than the observations at 2 m; this effect could also be explained by incorrect soil moisture initialization. The biases in α are mostly positive, i.e. the values of α from the forecasts are larger than those diagnosed from the observations. Because of the differences in height (10 and 60 m in the observations versus 14 and 53 m in the simulations) and the higher surface roughness length in the model, it is expected that the value of α is slightly larger in the forecasts than the observed value under neutral conditions.

Table 3.2: Error statistics between WRF forecast sets and the observations at Høvsøre for the period 1–30 October 2009. Only forecasts with lead times between 7–30 hours are considered. The levels used in the wind statistics for the observations and model are: 10 and 14 m, 60 and 56 m, and 160 and 164 m. The lowest BIAS or RMSE error per PBL experiment is in bold.

BIAS (WRF-observations) Experiment	ACM2	MRF	YSU	MYJ	MYNN2	MYNN3	QNSE
wind speed 10 m ($\mathrm{m\,s^{-1}}$)	0.68	0.32	1.19	**0.13**	0.34	0.59	-0.24
wind speed 60 m ($\mathrm{m\,s^{-1}}$)	**0.16**	0.34	0.37	-0.20	0.25	0.37	-0.33
wind speed 160 m ($\mathrm{m\,s^{-1}}$)	0.12	0.07	-0.27	0.04	**0.00**	0.08	0.23
Temperature 2 m (°C)	0.68	0.93	1.21	**0.22**	0.40	0.44	0.29
Temperature 100m (°C)	**-0.03**	0.32	0.20	-0.56	-0.28	-0.31	-0.65
α 10/60 m ($\times 10^{-1}$)	-0.09	0.33	-0.59	0.38	0.42	**0.22**	0.89

RMSE (WRF-observations) Experiment	ACM2	MRF	YSU	MYJ	MYNN2	MYNN3	QNSE
wind speed 10 m ($\mathrm{m\,s^{-1}}$)	2.02	2.15	2.43	**1.95**	2.06	2.06	2.06
wind speed 60 m ($\mathrm{m\,s^{-1}}$)	2.52	2.69	2.63	**2.49**	2.66	2.57	2.56
wind speed 160 m ($\mathrm{m\,s^{-1}}$)	2.49	2.67	2.53	**2.46**	2.56	2.47	2.65
Temperature 2 m (°C)	1.50	1.55	1.78	1.26	**1.24**	1.28	1.30
Temperature 100 m (°C)	1.32	1.42	**1.24**	1.43	1.38	1.43	1.50
α 10/60 m ($\times 10^{-1}$)	**1.30**	1.85	1.83	1.65	1.79	1.77	2.17

Very small differences in RMSE of wind speeds are seen among the various sets of forecasts, and the spread among them becomes even smaller with increasing height. The RMSE in temperature is also similar among the sets of forecasts at 100 m, but quite different at 2 m. The RMSE in α varies between 0.130 for the ACM2-based forecasts and 0.217 for the QNSE-based forecasts.

3.3.3.1   Diurnal cycle

Figure 3.7 compares the diurnal variations in the observations at Høvsøre with those predicted by the seven experiments in Table 3.1. Most experiments overestimated the wind speed at 10 m, especially after sunset (16:00–22:00 UTC). The overall lowest bias (0.13 $\mathrm{m\,s^{-1}}$) is that of the MYJ simulation; the highest that of the YSU simulation (1.19 $\mathrm{m\,s^{-1}}$). At 60 m, the wind speeds simulated by the various PBL experiments are within 1 $\mathrm{m\,s^{-1}}$ of the observations and most tend to underestimate its value during 1:00–7:00 UTC and overestimate it during 11:00–16:00 UTC. At 160 m, the spread in the wind speed among the various PBL experiments is much reduced from that at 10 and 60 m, but all the experiments fail to reproduce the diurnal range in wind speed in the observations (~ 2 $\mathrm{m\,s^{-1}}$).

All the forecast sets overestimated temperatures at 2 m (by up to 2 K) compared to observations; the YSU-based forecasts being the warmest and the MYJ-based forecasts the coldest. Biases are generally smaller during the daytime than during the nighttime. At the third model level ($\sim$105 m), the simulated temperatures are much closer to those observed at 100 m than modelled 2-m temperatures are to observed 2-m ones. The forecasts using TKE-based PBL schemes underestimate the temperatures at this level (100 m) while the non-TKE-based forecasts overestimate them. The ACM2-based forecasts are nearly perfect

Figure 3.7: Comparison of: (a) wind speed (m s^{-1}) at 10 m and first model level ($\sim$14 m), (b) wind speed (m s^{-1}) at 60 m and second model level ($\sim$54 m), (c) wind speed (m s^{-1}) at 160 m and fourth model level ($\sim$164 m), (d) heat flux (W m^{-2}) at 10 m, (e) temperature (K) at 2 m; (f) temperature (K) at 100 m and third model level ($\sim$105 m); (g) friction velocity (u_*; m s^{-1}) at 10 m and (h) α as a function of time of the day (UTC hours). Model values are extracted from the land grid point closest to Høvsøre (Figure 3.6a) for each of the experiments for the period 1–30 October 2009. Only 7–30 hours of forecasts are averaged.

with an overall bias of only 0.03 K. The phase of the diurnal cycle simulated at the third model level agrees well with the observations, whereas the phase of that simulated at 2 m seems to be 1 hour ahead of that observed. Consistently, the simulated surface heat flux shows the same lag as the 2-m temperature and is under-predicted at night and over-predicts the afternoon maximum by $\sim$0.3 W m^{-2} for the QNSE- and YSU-based forecasts. The spread among the forecast sets is small in terms of heat flux.

The value and diurnal evolution of the friction velocity is poorly represented in all the forecast sets. The spread among the forecast sets is smaller than their averaged bias to the observations. The morning and evening transitions in u_* in the observations, at about 9:00 and 15:00 UTC, exist at approximately the right time in the simulations, but the magnitudes of change of u_* in the course of the day are larger in all the forecast sets than in the observations. The values of u_* for the YSU-based forecasts have the highest biases of all the sets, but show very little variation during the day in better accordance with the observations.

The relationship between the values of α computed from observations and those in the forecast sets has an interesting pattern through the day. During the daytime (9:00–16:00 UTC), most sets follow quite closely the time evolution of the observed values. In the morning and at night, the spread among the experiments is much larger than during the day. The two outliers are the QNSE- and YSU-based forecasts: The first over-forecast the general evolution of α during the day, while the second ones show a nearly constant value throughout the day. A value close to that is expected from MOST under neutral stability conditions.

The relationship between the biases in 2-m temperature and in the temperature at the second model level is consistent with the errors in wind shear. While most simulated 2-m temperatures are warmer than observed, the temperatures at the second model level are under-forecast. Therefore, the thermal stratification in the forecast sets is lower than that observed (so it is closer to MOST neutral conditions).

3.3.3.2   Wind shear exponent

The value of α is examined as a way of diagnosing the wind shear, which is an important parameter in wind energy. The focus is on α between 10–60 m to avoid possible contamination with the internal boundary resulting from the sea-land transition at Høvsøre. As shown in Floors et al. (2011), from different IBL models and observations, the kink in the wind profile at Høvsøre occurs at 60–70 m. Also, this is the layer where the wind shear varies the most due to the presence of the Earth's surface.

Figure 3.8 shows the evolution of α during the period 1–30 October 2009 for all forecast sets and the observations. For every forecast set, $\alpha > 0.15$ dominate during the night and morning hours, whereas $\alpha < 0.15$ prevail during daytime. The distribution of α is fairly well captured by all the forecast sets, except for the YSU-based ones, which simulate neutral conditions for most of the month at all times of the day. The block of $\alpha > 0.20$ nighttime conditions during 9–19 October was captured best by the forecasts using the TKE-based parameterizations MYNN2 and MYNN3. The QNSE- and MYJ-based sets forecast too large values of α during that period. However, the MYNN2- and MYNN3-based sets showed larger values for the period from about 19 to 26 October than those derived from the observations. During
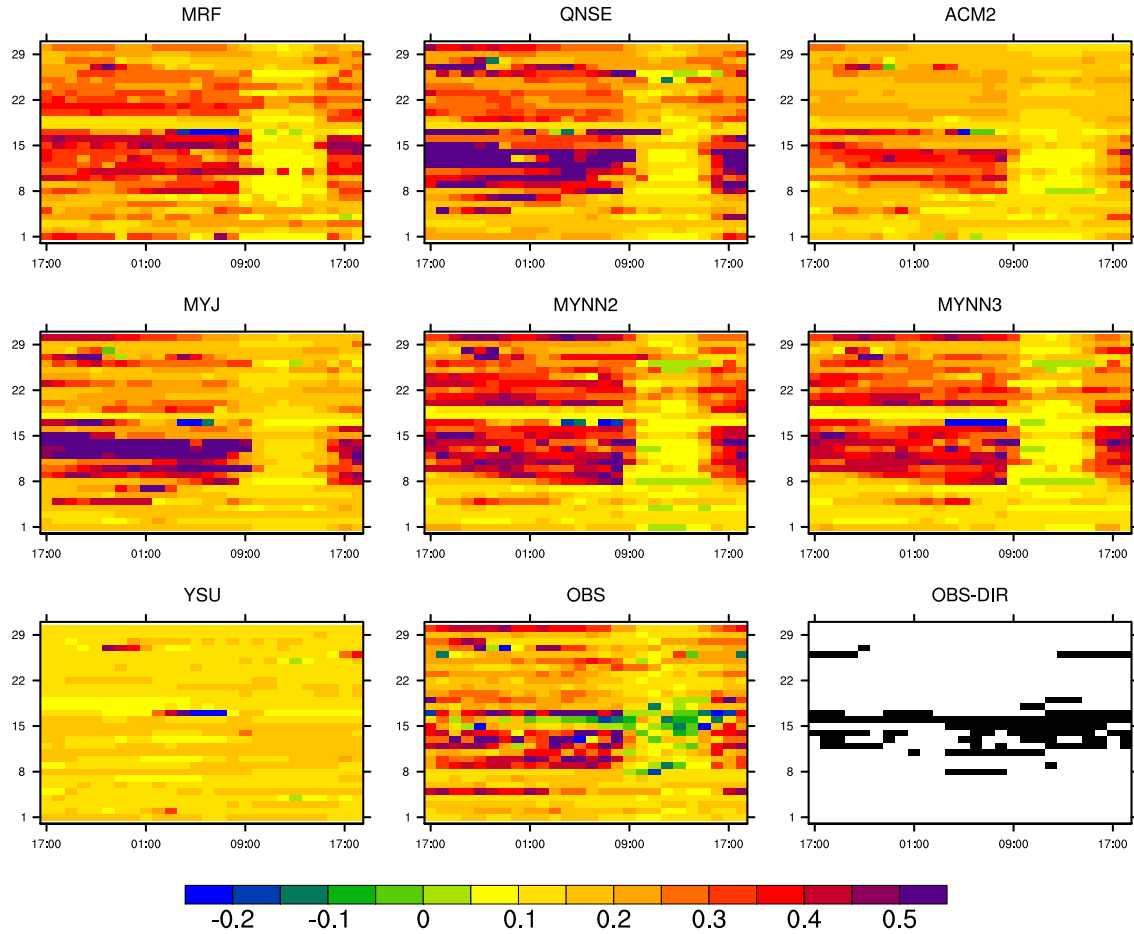
Figure 3.8: Comparison of the wind shear parameter α for Høvsøre as a function of time of the day (UTC, x-axis) and day of the month (y-axis) computed from the wind speed values at 14 and 54 m in the experiments and 10 and 60 m in the observations. The title above each subplot indicates the experiment name (Table 3.1). The bottom middle panel represents the observations. The bottom right panel displays by black squares the hours when the observed wind directions correspond to those in the wake of the wind turbines (330–30°).

this period the QNSE- and ACM2-based forecasts show better agreement with observations. On the other hand, the MYNN2- and MYNN3-based sets did capture the large values of α that occur during the last day of the month better than any other forecast sets. The very low observed values of α during the daytime are not well captured by any of the sets. However, some of the negative values observed during the period 10–18 October might be influenced by the wake of the wind turbines (last panel in Figure 3.8), which for these heights will tend to produce a more mixed boundary layer and thus lower values of α.

3.3.3.3 Curvature of the wind speed profile

The representation of the modeled α-parameter depends on how well the models prognose the wind speed profile, but due to the selected α levels the wind shear analysis above is valid for the layer close to the surface only. Hence, to examine the wind speed profile further up, averaged wind profiles for 1–30 October 2009 were calculated (Figure 3.9). The

Table 3.3: Boundaries of stability classes based on the observed Obukhov length. The classes were used in the grouping of wind profiles and computations of error metrics. Percentages of wake-free data with respect to all the data per stability class are also shown.

Stability class	Obukhov length [m] range	Number of samples (Percentage of wake-free data)
unstable	$-500 \leq L < -50$	55 (72%)
neutral	$|L| > 500$	200 (93%)
near stable	$200 < L \leq 500$	118 (93%)
stable	$50 < L \leq 200$	136 (83%)
very stable	$10 < L \leq 50$	72 (74%)

averages were separated into stability classes according to the observed Obukhov length, as indicated in Table 3.3. This analysis is based on the work of Gryning et al. (2007) and Peña et al. (2010) at Høvsøre. They showed that when plotting the wind speed normalized by the surface friction velocity as a function of the logarithm of height, the "stability" of each profile is easily distinguishable: a more or less straight, curving downward, and slightly curving upward line for neutral, stable, and unstable conditions, respectively. This behavior is expected from similarity theory and is easily seen in the observed wind profiles (black solid lines) in Figure 3.9. Because the mean surface friction velocity in the WRF simulations is always larger than that observed, there is a systematic offset between the observed and modeled normalized wind profiles. This is not relevant to the analysis presented in this section, since the focus is on comparing the shape of the profiles. The profiles in this figure include the wind speed averages at 10 m in the simulations, which is a diagnostic quantity of each PBL scheme. Note that a larger velocity range is used in the last panel to represent larger disparity among the PBL schemes under very stable conditions.

Figure 3.9 shows that the difference between observed and modeled normalized wind speeds increases with height in all stability classes, but especially for stability classes stable and very stable, because the winds simulated in most of the forecast sets are more typical of a neutral atmosphere. The forecast set that best captured the observed curvature of the average wind profiles is that QNSE based. The forecast set that had the most different average wind profile from the observed one is that YSU based; it shows the greatest deviation from the measured profile during stable conditions and produced "neutral" wind profiles in every stability class. The other sets lie in between these two with MYJ-, MYNN2-, MYNN3- and ACM2-based ones performing similarly. The largest spread among the sets is seen in the very stable class, but might not be as significant as in the other classes due to the reduced number of samples. The MYJ- and QNSE-based forecasts display higher winds at 10 m than at 14 m (Figure 3.9), which is a consequence of the diagnostic scheme used in these parameterizations.

3.3.3.4   Error metrics

When averaging wind speeds as in Figure 3.9, information about the daily performance of the experiments is included only if errors are systematic. However, in wind energy forecasting, one single poorly–forecast event can result in huge costs for transmission system operators. Therefore, a metric is needed that also includes information about outliers and penalizes them.
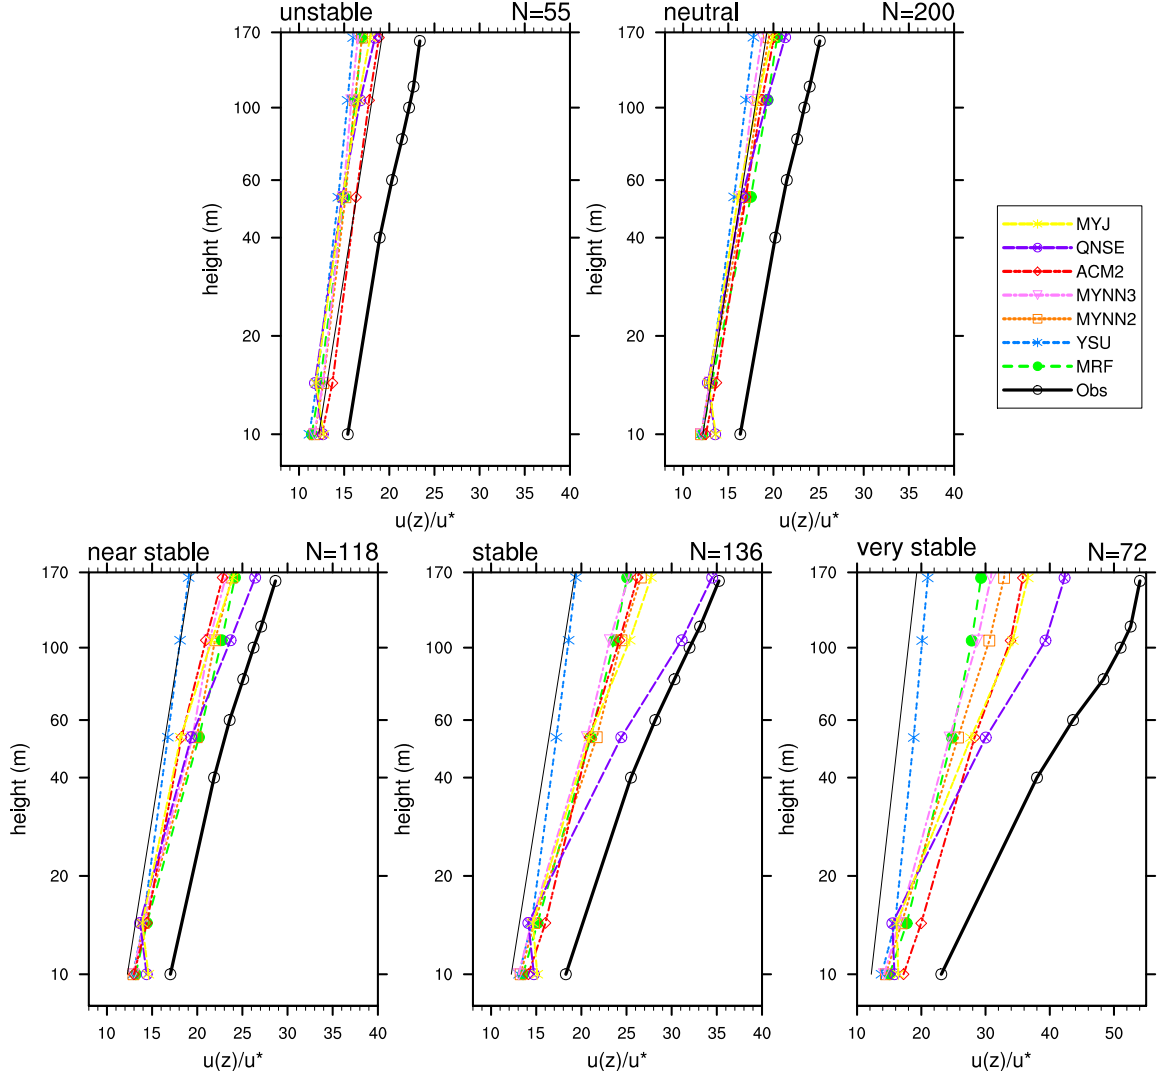
Figure 3.9: Comparison of the observed normalized wind speed profile (solid black thick line) and those simulated by the different PBL experiments for Høvsøre. The profiles are grouped into five stability classes as described in Table 3.3 and indicated in the title of the plots. The thin black solid line serves as a reference for a neutral profile with $z_0 = 0.05$ m. The number of observed cases is indicated by N in each plot. Times with wind directions from 330–30° are excluded from the calculations.

The RMSE was calculated for wind speeds at heights 60 and 100 m, for the α-value, and the WPE for every stability class in Table 3.4 to find objectively the best performing parameterization for wind profile forecasting.

The error metrics show that the performance of each experiment depends on the analyzed parameter, and in particular the observed atmospheric stability. For unstable conditions, the YSU experiment has the lowest RMSE for wind speed and the lowest WPE of all the experiments. For neutral and near-stable conditions, the ACM2 experiment is lowest in most wind speed and α RMSE. For stable and very stable conditions, the MYJ experiment is lowest in wind speed RMSE. Despite showing very good agreement with observations in forecasting the mean wind speed profile, the QNSE experiment, especially for classes unstable and neutral, has the highest wind speed RMSE and the largest WPE of all the experiments.

Table 3.4: RMSE for wind speed at 60 m and 100 m, for the wind profile, and the α-value for the seven PBL experiments in the five stability classes (Table 3.3) during observed wind directions between 30 and 330°. The lowest RMSE error per stability class and PBL experiment is in bold; the highest is underlined.

	ACM2	MRF	YSU	MYJ	MYNN2	MYNN3	QNSE
UNSTABLE							
wind 60 m (m s⁻¹)	2.212	2.372	**2.014**	2.212	2.229	2.264	<u>2.792</u>
wind 100 m (m s⁻¹)	2.409	2.556	**2.169**	2.394	2.377	2.410	<u>3.011</u>
WPE (m s⁻¹)	2.415	2.561	**2.194**	2.421	2.390	2.427	<u>3.008</u>
α	0.078	0.084	<u>0.107</u>	0.074	0.079	0.077	**0.069**
NEUTRAL							
wind 60 m (m s⁻¹)	**2.355**	2.384	2.449	2.579	2.587	2.380	<u>2.750</u>
wind 100 m (m s⁻¹)	2.377	2.389	2.462	2.452	2.501	**2.306**	<u>2.640</u>
WPE (m s⁻¹)	2.370	2.379	2.463	2.481	2.510	**2.321**	<u>2.676</u>
α	**0.048**	<u>0.107</u>	0.054	0.052	0.103	0.085	0.082
NEAR STABLE							
wind 60 m (m s⁻¹)	**1.996**	2.044	2.078	2.118	<u>2.321</u>	2.040	2.154
wind 100 m (m s⁻¹)	**2.048**	2.260	2.194	2.062	<u>2.288</u>	2.075	2.074
WPE (m s⁻¹)	**2.060**	2.221	2.198	2.121	<u>2.329</u>	2.111	2.171
α	**0.068**	0.124	0.074	0.101	<u>0.145</u>	0.131	0.143
STABLE							
wind 60 m (m s⁻¹)	2.102	<u>2.238</u>	1.944	**1.820**	2.127	2.035	1.922
wind 100 m (m s⁻¹)	2.237	<u>2.481</u>	2.019	**1.985**	2.160	2.102	2.117
WPE (m s⁻¹)	2.257	<u>2.462</u>	2.093	**2.011**	2.215	2.142	2.1647
α	**0.111**	0.118	0.143	0.141	0.142	0.131	<u>0.249</u>
VERY STABLE							
wind 60 m (m s⁻¹)	1.619	<u>1.753</u>	1.599	**1.446**	1.538	1.551	1.518
wind 100 m (m s⁻¹)	1.833	<u>2.110</u>	2.002	**1.453**	1.643	1.619	1.590
WPE (m s⁻¹)	1.811	<u>2.060</u>	1.951	**1.531**	1.685	1.705	1.671
α	**0.172**	0.220	<u>0.310</u>	0.181	0.210	0.207	0.233

3.3.3.5   Effect of land surface and surface layer schemes

Figures 3.7 and 3.8 showed that the results from the forecast set using the YSU parameterization gives an almost constant value of α throughout the day and a wind profile typical of neutral conditions most of the time, whereas the other sets show variations of the α-parameter more similar to those observed. The obvious question is: Are the YSU-based forecasts, and to a lesser degree all the other forecasts sets, deficient in the response of the wind profile to the stability conditions or are these conditions already misdiagnosed in the simulations? To answer this question, Figures 3.10 and 3.11 show the relationship in surface heat flux, surface friction velocity, Obukhov length, and α between the observations and the forecast sets. Figure 3.10 is for the forecasts using the TKE-based parameterizations, while Figure 3.11 is for the first order and non-local-K based PBL parameterizations. All the forecast sets use the Noah land surface model (LSM), except for ACM2 that uses the Pleim-
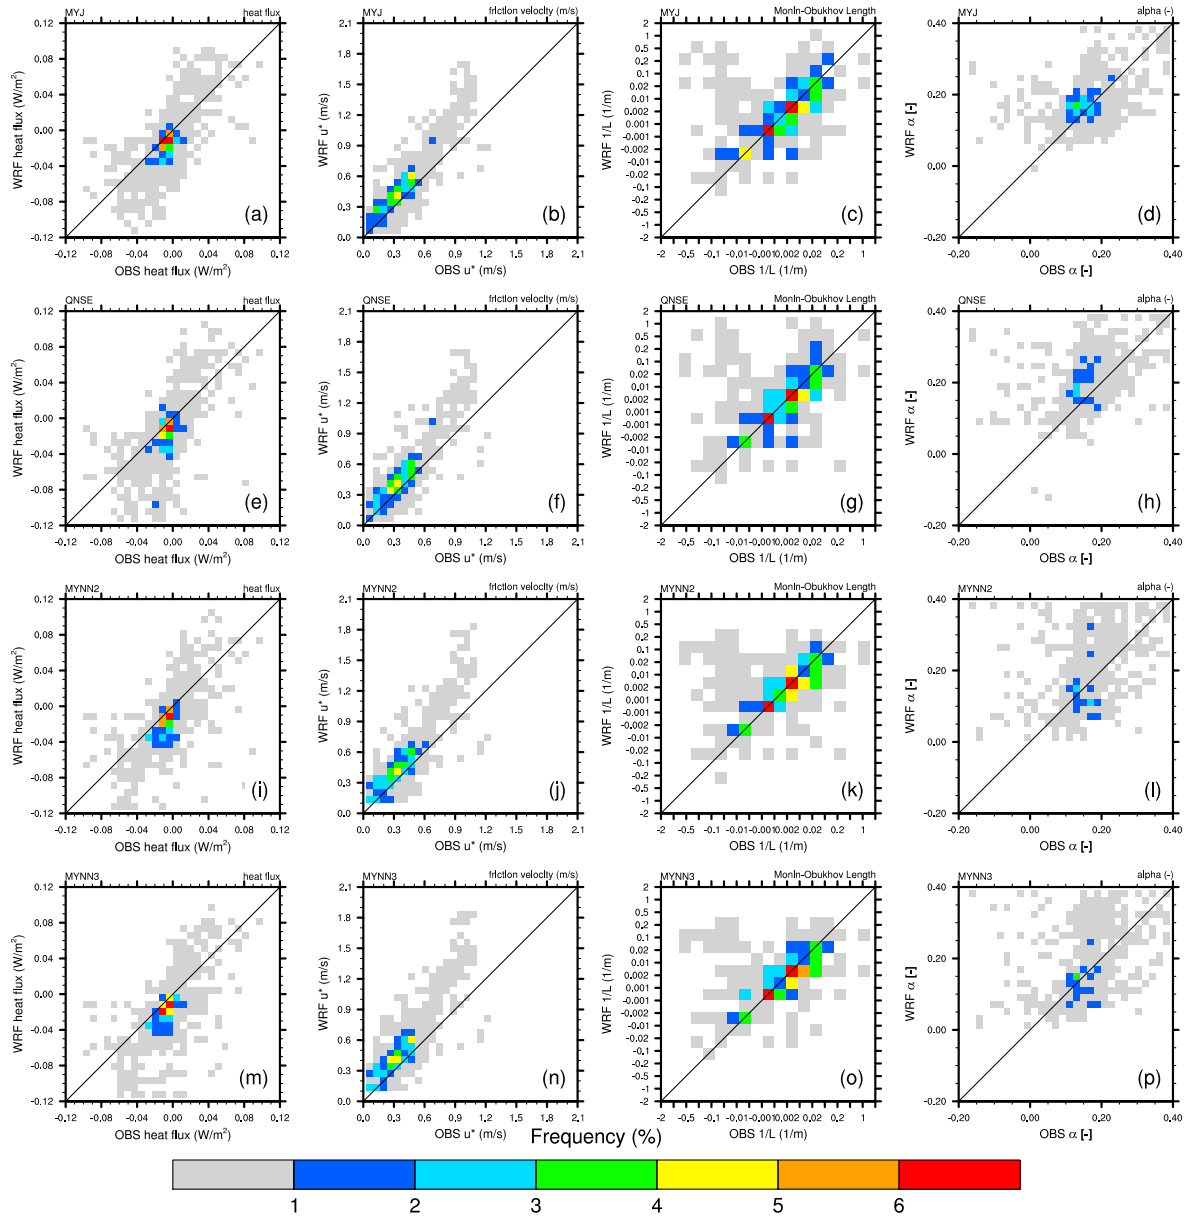
Figure 3.10: Comparison of observed (x-axis) versus WRF-simulated (y-axis) fields for the experiments using TKE-based PBL parameterizations: (a)–(d) MYJ, (e)–(h) QNSE, (i)–(l) MYNN2, (m)–(p) MYNN3. The four columns from left to right show: heat flux (W m^{-2}), surface friction velocity (m s^{-1}), Obukhov Length ($1/L$) and α-parameter. Values are at Høvsøre during 1–30 October 2009. The label bar shows the frequency of occurrence (%) of each 2D bin. Times with observed wind directions from 330–30° are excluded from the calculations.

Xu LSM and a similar surface layer based on similarity theory, but with slightly different parameters and stability functions.

The simulated values of heat flux in the seven forecast sets are overall lower than those observed. But the range of heat flux in the simulations is larger than that observed. The relationship between observed and simulated friction velocity is reversed: the forecasts over-
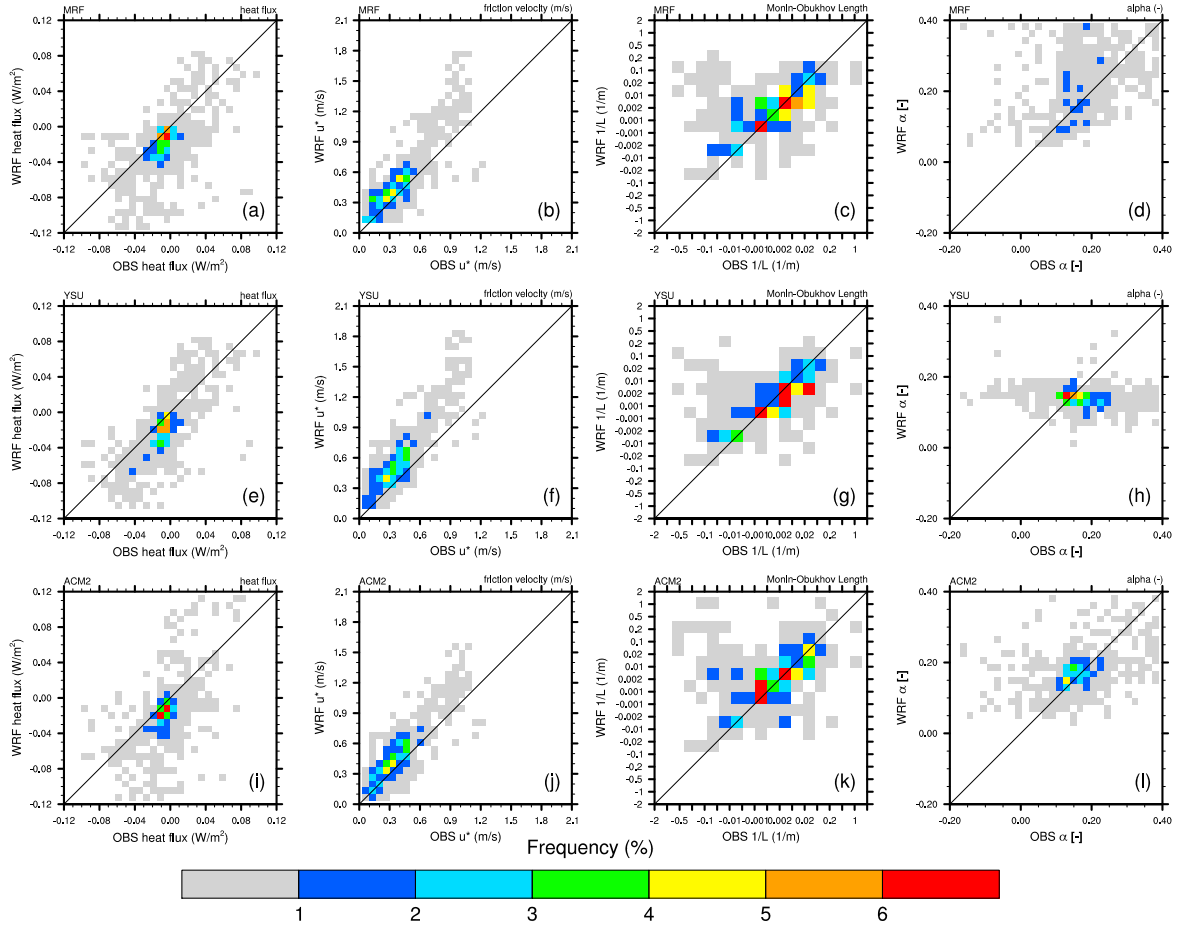
Figure 3.11: Same as Figure 3.10, but for the experiments using the first order and non-local-K based PBL parameterizations: (a)–(d) MRF, (e)–(h) YSU, and (i)–(l) ACM2.

estimate u_*. In the MYNN2-, MYNN3-, MRF- and YSU-based forecasts the value of u_* is constrained to a minimum value of 0.1 by the PBL parameterization implementation. Because L is calculated as the ratio of these two previous quantities, their errors tend to compensate each other resulting in a good relationship between their observed and simulated values, especially for the simulations using the TKE-based PBL schemes (Figure 3.10). Reinforcing the results from Figure 3.7, the relationship between observed and simulated heat flux and surface friction velocity plots in Figures 3.10 and Figure 3.11 is very similar for all the forecast sets. In terms of $1/L$, the sets using the TKE-based PBL schemes show higher correlations between observed and simulated values than those using the other PBL schemes. In terms of α, the relationship between observations and simulations varies greatly from simulation to simulation showing considerable spread. The highest correlations between observed and simulated α-values are in the ACM2 ($r = 0.51$) and MYJ ($r = 0.42$) forecasts; the lowest in the YSU-based ones ($r = -0.16$).
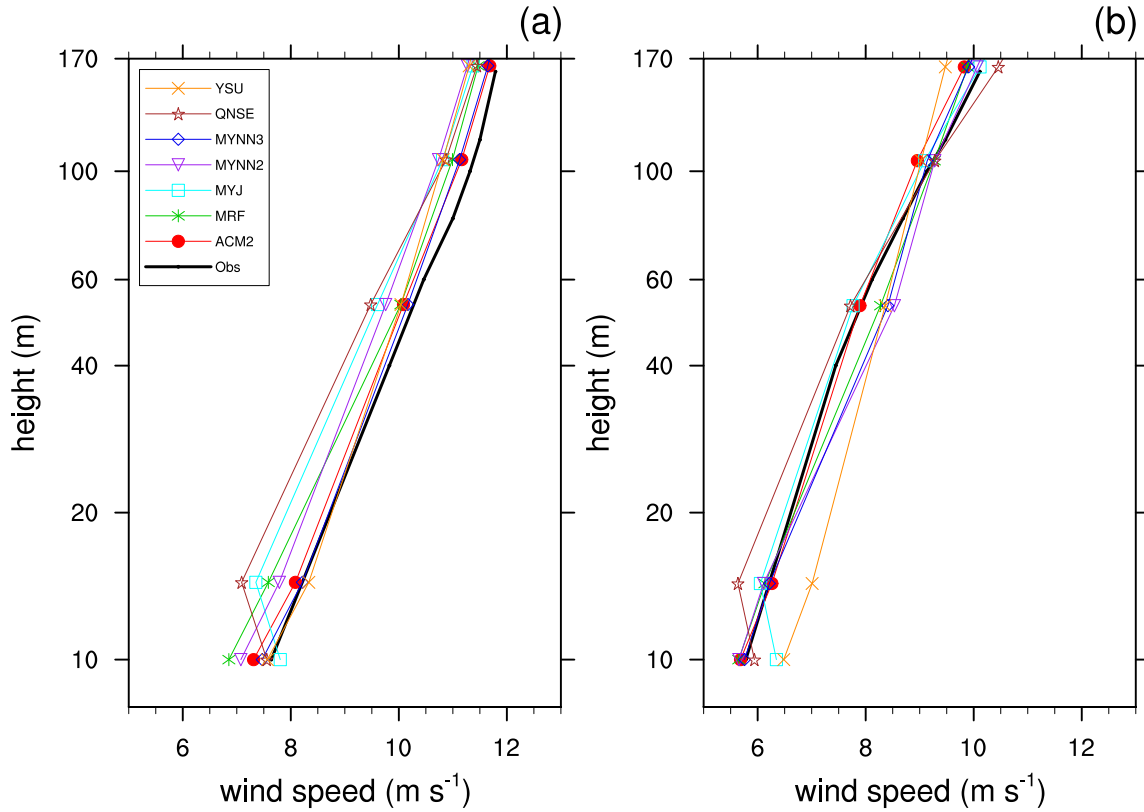
Figure 3.12: Time-averaged wind speed (m s^{-1}) as a function of height for observations and forecasts at Høvsøre during the periods: (a) 1–14 October and (b) 15–30 October 2009. Only values from forecast lead times of 7–30 hours are used.

Figure 3.13: Wind roses for 60 m wind observations at Høvsøre during the periods: (left) 1–14 October and (right) 15–30 October 2009. The mean wind speed for each wind direction is given.

3.3.4 Discussion

Boundary layer winds simulated by the WRF model using seven different PBL schemes have been evaluated against observations at a coastal site in Denmark with a focus on the representation of the wind shear under different atmospheric conditions. The results show

that although the time-averaged wind speeds predicted by a NWP model agree well with those observed, the details in the simulated variations in time and in height, which are important in wind energy applications, are much less in agreement with observations.

Taking into account all the metrics in Tables 3.2 and 3.4, the best performing forecast set for this site and this period is the MYJ-based set. It shows the lowest WPE under stable and very stable conditions. Under other stability conditions its WPE is usually very close to the lowest value. While the QNSE-based forecasts produced average wind profiles with a curvature similar to that of the observations, its RMSE values (Table 3.4) show that the individual profiles deviated considerably from those observed. Similarly, the probability distribution of α (Figure 3.10h) is biased towards higher values.

The results seen in the verification of the simulations reported here are in line with previous studies. Storm and Basu (2010) found that simulations using the YSU parameterization produce "too neutral" wind profiles. They explain its failure by noting that the parameterization has excessive mixing in the stable PBL, which destroys the near-surface shear. Shin et al. (2011) also mentioned excessive mixing in the ACM2 and YSU scheme. They further explained that in all their simulations for a single day, the value of u_* is overestimated and the schemes are incapable of reproducing the decoupling of the air at the surface and aloft at night and thus the occurrence of low-level jets. This study shows the same problem: too high winds occur near the surface at night (Figure 3.7a) except for in the QNSE scheme, which is the only one trying to "bend back" the wind profile above 100 m as seen in the observations (Figure 3.9; very stable). Hu et al. (2010) concluded that the differences between their evaluated PBL schemes ACM2, YSU and MYJ were due to differences in vertical mixing strength and entrainment of air from above the PBL. The latter was not evaluated in this study, since I focused on low-level boundary layer winds available from the tower measurements at Høvsøre. Verification of winds from WRF simulations against lidar measurements up to 2 km is currently underway at Høvsøre.

The observed temporal variations in wind shear clearly indicate that using 10 m winds alone is not sufficient when verifying mesoscale model simulations for wind energy applications. Firstly, an accurate simulation of winds at 10 m does not guarantee an accurate simulation of wind conditions at hub height, nor across the whole rotor area. It has also been shown in Figure 3.9, that the 10-meter winds diagnosed by the QNSE and MYJ parameterizations are higher than those that could be expected from similarity theory (Jimmy Dudhia, personal communication). Secondly, the wind field near the ground (e.g., 10 m) is greatly affected by the local topography, roughness, and the presence of obstacles. Therefore, the representativeness error of wind measurements is larger than that from other atmospheric variables. Because of the inherent design and limited resolution of most NWP mesoscale models, such effects are not adequately represented in their simulations. It is thus questionable to directly compare the raw wind output of a mesoscale model with site observations at heights lower than 40 m, unless the site and model surroundings are very flat and homogeneous or it is located offshore. This fact is partly demonstrated by the larger bias in the 10-meter wind than those at 60 or 160 m. A post-processing technique that uses the output from a micro-scale simulation of the site conditions using the Wind Atlas Analysis and Application Program (WAsP; e.g., Troen and Petersen (1989)) can be applied to the mesoscale-model derived winds to compensate for the features not taken into account by the NWP model as done in Larsén et al. (2010). This technique was not applied to the results presented here because the terrain is very flat and homogeneous, except for the coastline.

Since the choice of PBL scheme in WRF is tied to a certain land surface and surface layer scheme, the results of this study are not only related to the choice of PBL scheme. Shin and Hong (2011) found that the sensible heat flux depends on the surface layer parameterization rather than on the temperature gradient near the surface (i.e., on the PBL scheme). They stated that the value of the forecast 2-m temperature depends on the LSM, but that the wind and temperature profiles are determined by the PBL mixing algorithms. In addition, the difference in temperature bias between 2 and 100 m can be explained by inadequate grid spacing to resolve the transition zone between the coast and the nearest grid point to the measurement mast on land. If the flow is from western directions, which are prevailing in this climate, air parcels moving from the ocean to land have no chance to cool as they are advected over the adjacent grid point used for verification. Land-surface variability impacts turbulent surface fluxes. The variability of surface fluxes is correlated to roughness length, atmospheric stability and advection (Talbot et al., 2010). A more accurate roughness length representation could therefore result in better forecasts. Indeed, the roughness length in the WRF model changes on 15 October from 0.15 to 0.05 m, which is in the middle of the evaluation period. Figure (3.12) shows the time-averaged wind speed as a function of height for the observations and the simulations for the periods. Figure 3.13 shows the wind roses for these two periods derived from the observations at 60 m. The first half of October 2009 is dominated by strong offshore winds, while during the second half easterly winds dominate. Therefore, mean wind speeds are stronger during the first half, when surface roughness lengths are larger, than during the second half when surface roughness lengths are smaller. The effect of the change in roughness on 15 October is thus minimized because of these opposing changes. Very little sensitivity is found from additional simulations modifying the WRF surface roughness to that observed at the site (not shown). This is partly due to the strong horizontal diffusion in WRF, which makes the grid point chosen for verification more like ocean than land, especially when the wind comes from the sea, which is common at the site.

The value of z_0 used in the simulations affects the value of the friction velocity. Since the value of the roughness length specified in the WRF LSM is larger than that observed at Høvsøre, the model physics compensate the value of other parameters to produce wind forecasts closer to reality. This is done by artificially enhancing the value of u_* (Figure 3.7g; second column of Figures 3.10, and 3.11). Changing z_0 in the model to more realistic values could therefore improve the simulations of u_*, but might result in unrealistic values of the simulated surface sensible and latent heat fluxes.

The Obukhov length is fairly well predicted by all the seven PBL schemes (last column of Figures 3.10 and 3.11). Furthermore, these Figures show that the heat flux, u_* and L are fairly well correlated in most of the schemes, but this is not the case for α. Separating the wind profiles in classes according to observed Obukhov length reveals that the PBL schemes tend to diffuse the surface momentum fluxes upward in a way that is not always consistent with the MOST. This is partly because the surface momentum and heat fluxes are not realistically predicted, although their combination produces relatively reasonable L values. Similar results are found in Peña and Hahmann (2011).

It can be argued that some of the problems in the representation of the wind shear in the WRF simulations could be due to the relative low vertical resolution of the model grid. Additional simulations were carried out with enhanced vertical resolution (63 vertical levels; not shown). The obtained results in terms of the shape of the wind profile are almost identical to those presented here. However, this conclusion might not hold over more complex terrain.

3.4   Summary

This study analyzes the results of numerical simulations with the WRF model using seven different PBL parameterizations for their performance in wind energy forecasting. A verification methodology was further presented for the use of mesoscale models for wind energy that takes the vertical structure of the wind into account. Verification metrics include averaged time series, the α-parameter (a measure of the vertical wind shear), averaged wind profiles and RMSE, which were grouped into atmospheric stability classes. The results clearly indicate that for wind energy applications, i.e., wind power forecasting and wind resource assessment, verification against 10-m wind speeds alone is not sufficient. Validating a NWP model at different heights will lead to more accurate guidance. This becomes even more important in the future, as wind turbines get taller.

The ability of the seven forecast sets to simulate the mean wind speed and its time variability depends strongly on atmospheric static stability. None of the PBL schemes is able to outperform the others under the range of stability conditions experienced at Høvsøre. The MYJ parameterization performs best during stable and very stable atmospheric conditions and the ACM2 during neutral and near stable cases. The YSU scheme outperforms the others during unstable conditions. Average wind profile shapes of the QNSE scheme compare best of all schemes with the average observed wind profiles, but individual profiles actually differ the most from those observed. The YSU PBL parameterization does not exhibit correct diurnal variations and curvature of the average wind profiles; it produces wind shears typical of the neutral atmosphere most of the time due to strong vertical mixing near the surface during the nighttime (Shin and Hong, 2011; Hu et al., 2010).

The choice of the best model setup for a forecasting system for a particular region will thus depend on the typical distribution of atmospheric stability conditions at the site. For regions where stable conditions prevail, the MYJ PBL scheme is recommended. For regions where neutral and near stable cases dominate, the ACM2 PBL scheme is recommended. Thus, an operational wind power forecasting system giving the best forecast under all stability conditions could consist of a multi-scheme ensemble with weights depending on the stability. These conclusions are of course not universally applicable, because the verification presented in this study focuses on a very flat and homogenous land site. Similar conclusions were drawn for nearly the same WRF simulations over the North and Baltic Seas in Hahmann and Peña (2010). The performance of the various schemes over complex terrain is not possible to determine from the results presented here. In addition, for an evaluation of results from NWP models in more complex terrain, the representativeness of measurements and model output would have to be taken into account.

The analysis shows that the ability of the seven PBL schemes to capture the shape of the wind profile also depends on atmospheric stability. The curvature of the wind profile simulated by all the schemes diverted more from the observations the more stable the atmosphere. This deviation increases at 100 and 160 m, because of the schemes' tendency to produce profiles that are more mixed and thus with less vertical shear than those observed. An overestimation of u_* and 2-m temperature, and an underestimation of the second model level temperature were found to be partly the reason for the schemes to produce wind profiles with vertical shear expected from a neutral atmosphere. All the schemes used in the simulations underestimated the wind around turbine hub height during the night and overestimated it during the day. This diurnal compensation is coincidentally good for wind energy resource assessment based on NWP model output at 60–100 m height. It is also true that the wind speed simulated by

the NWP model will somewhere in the atmosphere match the observations. This will not happen close to the surface due to the over-prediction of u_* but somewhere higher from the ground, within the height of common wind turbines. This is important for accurate wind power forecasting, as it will pay off to find the best performing level at each site.

In this study the focus was on simulating wind and its vertical structure using seven WRF PBL parameterizations. Analyses concerning model behavior in predicting eddy diffusivity, TKE, Prandtl number, mixing length, virtual potential temperature profiles and boundary layer heights can lead to more explanations on the different behavior of the schemes. This would be a suggestion for further studies and was partly conducted for the YSU, ACM2, MYJ, and QNSE PBL parameterizations in Shin and Hong (2011). Furthermore, since a PBL parameterization is not applied in isolation from the other settings of the model, interesting future work would be to explore other aspects of the model within the verification framework that is proposed in this paper. Evaluating the forecasts at different locations could lead to interesting results as well. A big shortcoming is the lack of tall measurement masts and the low availability of flux measurements at most sites. A longer period of study would have provided more solid conclusions; however, additional simulations at Høvsøre and other sites have shown that the conclusions are robust.

3.5   Conclusions

This chapter examined mesoscale forecasts for wind energy purposes and presented the results of experiments using the WRF model. A first assessment of hub height wind forecasts demonstrated the need for more detailed analyses of wind speeds and was the motivation for a more extensive study, in which seven PBL schemes of the WRF model were evaluated for wind energy purposes. This study stated the importance of wind *profile* verification when it comes to wind energy forecasting. Furthermore, the ability of the seven experiments to forecast the mean wind speed and the wind shear and their time variability strongly depends on atmospheric static stability. Wind forecasts from the experiment using the Yonsei University PBL scheme compared best with observations during unstable atmospheric conditions, while the Asymmetric Convective Model version 2 PBL scheme did so during near stable and neutral conditions, and the Mellor-Yamada-Janjic PBL scheme prevailed during stable and very stable conditions.

Mesoscale models have been widely used for forecasting the wind at hub height and a large efford towards the aim of meteorological forecasts for wind energy purposes has been made by the ANEMOS project (Giebel et al., 2011). As discussed in detail in section 3.3.1, the WRF model has been evaluated previously for wind energy purposes. However, the research paper included in this thesis (Draxl et al., 2011b) is a more extensive comparison of seven PBL parameterizations available in WRF than has been done before and takes into account several heights across the rotor area.

A verification methodology for mesoscale models is presented in this paper, that focuses on the vertical variation of wind. Hopefully, it will convince others engaged in wind energy to refrain from using the 10 m winds only in wind energy forecasting.

CHAPTER 4

Improving Initial Conditions of Mesoscale Forecasts

Many aspects of a mesoscale model determine the quality of the forecast PBL structure (Draxl et al., 2011b), and thus the forecast winds. One of these aspects is the accuracy of the initial conditions. Many studies suggest that improved initial conditions have the potential to improve the subsequent forecasts (e.g., Hacker and Snyder (2005), Rabier (2005). Liu et al. (2006)). The analysis of the Poor Man's ensemble in Chapter 3.2 further suggested that the error statistics were better for the experiments that used initial conditions where more observations were assimilated. This chapter deals thus with introducing the idea of data assimilation to improve short-term wind energy forecasts.

4.1 Data assimilation

The basic concept of data assimilation is to produce an accurate image of the true state of the atmosphere at a given time. This accurate image is referred to as the *analysis*. The analysis can then be used as input or initial state for a numerical weather forecast. Due to the nonlinearity of the governing equations in a NWP model, errors in the initial conditions grow in time. Having an analysis as accurate as possible is therefore the key motivation in data assimilation. Data assimilation is about finding the model state most consistent with the observations, taking advantage of consistency constraints with laws of time evolution and physical properties.

The basic problem in data assimilation is to combine different sources of information: information about observations, about the NWP model, about physical laws and dynamics, and the error characteristics of the observations and the model. The idea behind combining this information is the Bayes' theorem.

The Bayes' Theorem states how the (prior) statistical knowledge is updated in the light of new information:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}, \tag{4.1}$$

where $P(A)$ is the prior probability of event A knowing nothing about event B, and $P(A|B)$ is the posterior probability of event A given that event B is known to have occurred. With the events defined as O (the observations), B (the model background or a priori estimate of the model state) and A (the analysis as the most probable state of the system), the aim is to find the most likely analysis given the observations O and the background B:

$$P(A|O \text{ and } B) = \frac{P(O \text{ and } B|A) \cdot P(A)}{P(O \text{ and } B)} \tag{4.2}$$

The information from the observations and the background are associated with covariance matrices corresponding to their respective errors (Pannekoucke et al., 2008). These error covariance matrices determine the respective weights given to each piece of information in the analysis, e.g., they determine whether the observation should have more weight than the background, or vice versa. More weight should be given to the background, if the model performs generally well, i.e., the model is likely to produce an accurate forecast. The observations should be given more weight, if they are known to be of good quality; then it is desired that the observations influence the model solution to a greater extent. The correct specification of these statistics remains a challenge in data assimilation systems.

Many assimilation techiques have been developed for NWP. Assuming Gaussian probability density functions in the Bayes' Theorem, the main techniques in atmospheric data assimilation widely used operationally are

- Four Dimensional Data Assimilation

- Variational data assimilation methods

- Ensemble Kalman Filters

4.1.1   Four-Dimensional Data Assimilation

Four-Dimensional Data Assimilation (FDDA), is based on Newtonian relaxation (i.e., nudging). FDDA belongs to continuous data assimilation techniques, which means that observations are assimilated at the time that they occur. This is a data assimilation method that relaxes the model state toward the observed state by adding tendency terms based on the difference between the two states to one or more of the prognostic (not diagnostic!) equations (Stauffer and Seaman, 1993). These terms force the model solution at each grid point towards observations which are distributed nonuniformly in space and time (observation nudging), or towards analyses of observations (analysis nudging)(Warner, 2011). Only observations that are model variables can be nudged. Indirect observations, like satellite-measured radiances, cannot be assimilated directly. These would have to be converted to model variables off-line, which can be a complex endeavor. The experiments in chapter 6

are performed with the approach of observation nudging. The algorithm uses only those observations that fall within a predetermined time window centered around the current model time step. The set of differences between the model and the observed state is computed at the observations locations. These "corrections" are then analyzed back to the grid within a region of influence surrounding the observations. This region is circular.

A prognostic equation for any dependent variable f, which applies at a particular grid point x and at a particular time step t, will thus look like

$$\frac{\partial f}{\partial t} = F(f,x,t) + \frac{f_{obs} - f}{\tau(f,x,t)} = F(f,x,t) + G(f)W(x,t)\epsilon(f,x)(f_{obs} - f) \qquad (4.3)$$

where F represents all the physical-process terms, f_{obs} is the observed value of f interpolated in three dimensions to the grid point, and τ is the relaxation time scale (Warner, 2011). This time scale is based on empirical considerations and is the time used to force the integration towards the observations. Multi-level observations are usually interpolated to model sigma levels. The relaxation-term weight can be separated into three components: the factor that determines the magnitude of the term relative to the physical terms in the equation (G), the function that defines the spatial and temporal influence of observations (W), and the observation-quality factor (ϵ), which is a function of data density. The horizontal weighting or influence is a Cressman function given by

$$w_{xy} = \begin{cases} \frac{R^2 - D^2}{R^2 + D^2} & D \leq R \\ 0 & D > R \end{cases} \qquad (4.4)$$

(Stauffer and Seaman, 1990), where D is the distance from the observation location to the grid point, and R is the horizontal radius of influence.

If the relaxation time scale is too small, the model solution will converge to the observation too quickly and the other variables will not have sufficient time to dynamically adjust. If the time scale is too large, errors in the model solution will not be corrected by the observations. FDDA seeks the minimization of the sum of the squares of errors between the model solutions and observations, distributed in space and time (Seaman, 2000). Optimization is attained when the model error is a minimum. Advantages of this method are that it is computationally efficient and robust, conceptually simple, and that it allows the model to ingest data continuously rather than intermittently (Reen and Stauffer, 2010; Stauffer and Seaman, 1990). The introduction of a small change to the model over multiple timesteps limits insertion noise and allows the model solution to achieve a greater degree of dynamic balance than intermittent methods (Reen and Stauffer, 2010; Seaman, 2000). The full model dynamics are part of the assimilation system, so that the analysis contains all locally forced mesoscale features.

Most operational data assimilation techniques assume that the errors, i.e., the weight with which the observations are assimilated, are isotropic (Otte et al., 2001). The data can thus be applied within circular regions of influence around measurement sites. Circular isotropic regions of influence are indiscriminate toward thermal and wind gradients that may reflect changes of air mass. The assumption of isotropic areas of influence can be especially erroneous in complex terrain, near coastlines or during synoptically forced conditions (like

frontal boundaries). More advanced data assimilation techniques like the Ensemble Kalman filter or Four-Dimensional variational data assimilation have the advantage of flow-dependent weights. Studies that investigate the benefit of anisotropic weighting of observations include, among others, Otte et al. (2001) or Benjamin and Seaman (1985). Ideally, in limited-area model data assimilation the radius of influence of an observation should depend on observation type and density, model error and resolution (Meng and Zhang, 2011).

Nudging has been used successfully over the last decades, first by Anthes (1974), and then by Stauffer and Seaman (1990) and Stauffer et al. (1991). The topics with nudging techniques over the last years include air quality modeling (e.g., Seaman (2000), Otte (2008)), surface data assimilation in connection with land surface models (Childs et al., 2006; Pleim and Gilliam, 2009; Pleim and Xiu, 2003) and the use of a weather analysis and forecasting system based on FDDA (Liu et al., 2008); Ruggiero et al. (1996), Alapaty et al. (2001) and Reen (2007) focused on nudging near-surface observations. Further applications of nudging include studies in complex terrain (Fast, 1995), of radiance measurements (Hayden, 1973), in the marine boundary layer (Leidner et al., 2001), the urban boundary layer (Liu et al., 2006) or using nudging to generate climatographies (Hahmann et al., 2009).

4.1.2   Variational data assimilation

Many operational numerical weather prediction centers use variational data assimilation methods to estimate the state of the atmosphere for weather prediction (Rabier, 2005). Variational data assimilation methods are also continuous, i.e., the data can affect the solution at every time step. They use an apriori (or background) state, which provides information otherwise missing from observations, and provide a realistic reference state needed to form the nonlinear observation operators used to assimilate many of the indirect observations (Bannister, 2008). In variational approaches both direct and indirect observations, like satellite-radiances, can be assimilated. The background state is error-prone and expressed by the background error covariance matrix, which describes how errors of different variables are correlated in space. It is derived from a short numerical forecast, so background errors share properties with forecast errors. Another way to say it is that the background error covariance matrix contains balance information in statistical form.

The evaluation of the error between the model solution and the observations is done through a cost function J (equation 4.5), which assumes that observation and background error covariances are described using Gaussian probability density functions (Barker et al., 2004). The optimal estimate of the true state of the atmosphere, i.e., the analysis, is determined by a minimization of the cost function through an iterative process (Seaman, 2000). Thereby it is assumed that the model can be neglected as a source of error during the assimilation period ("perfect-model" assumption). The cost function J is described by

$$J(\mathbf{x}) = \frac{1}{2}(\mathbf{x} - \mathbf{x}_b)^T \mathbf{B}^{-1}(\mathbf{x} - \mathbf{x}_b) + \frac{1}{2}[H(\mathbf{x}) - \mathbf{y}]^T \mathbf{R}^{-1}[H(\mathbf{x}) - \mathbf{y}]] \qquad (4.5)$$

where $\mathbf{x}$ is the state variable to be minimized, $\mathbf{x}_b$ the state of the model background, $\mathbf{y}$ the observation, $\mathbf{B}$ the background error covariance matrix, $\mathbf{R}$ the observation error covariance

matrix, and H the observation operator. Depending on the choice of $\mathbf{x}$ and H different variational data assimilation methods can be formulated. When $\mathbf{x}$ is considered as the state over the 3D spatial domain at analysis time, with H spatially interpolating this state and converting model variables to observed quantities, then it is called 3DVAR. When $\mathbf{x}$ is considered as the state over the 3D spatial domain *and* over the period for which observations are available, while H spatially *and temporally* interpolates this state and converts model variables to observed quantities, then this is called 4DVAR. Thus, 4DVAR determines the analysis state at every gridpoint and at every time within the analysis window i.e., a four-dimensional analysis of the available asynoptic data (Yannick Trémolet, 2009).

Compared to FDDA, the assimilation of data in a variational data assimilation system can give different results because the background error covariance matrix in a variational data assimilation system spreads out information in the vertical and horizontal directions in space, weights the importance of the a priori state and spreads information to other variables differently than in an FDDA system, where tendency terms are added to prognostic equations. The effect of the background error covariance in 3DVAR can be seen in Figure 4.1. where a pseudo observation for the horizontal u component of wind of $1 \, \mathrm{m \, s^{-1}}$ was assimilated. This is the outcome of a so-called single observation experiment, which is a way of diagnosing and tuning the background error covariance. The magnitude of the increments shown in Figure 4.1 is very small, but can be scaled depending on whether the observation or the background should have more weight in the data assimilation process. The equations of motion influence the way in which errors between and within variables are correlated (Bannister, 2008). The atmosphere is in a state of hydrostatic and geostrophic balance on most scales resolved by global models. Hence the probability density function that is represented by the background error covariance must allow only those likely background states that contain this near-balanced property. Consequently, if at a certain location the background state is perturbed through assimilation of a measured quantity, adjustments that are in near balance with the perturbation should be occuring elsewhere. This is why the assimilation of the pseudo u value leads to the depicted shapes in the v fields in Figure 4.1. The winds have a structure that is consistent with the adjustments due to geostrophic relations. Near the surface, where the pseudo observation was assimilated (in this example on model level 5), the effect of the coriolis force can be seen by rotated analysis increments.

An example of the assimilation of only wind speed observations into FDDA and 3DVAR is shown in Figure 4.2. The increments show a circular radius of influence for wind in the case of FDDA, where balance constraints are visible as off-circular increments. The increments of 3DVAR show a well defined dipol-structure. Important in these two figures discussed here is not the magnitude of the increments, which can be tuned, but their structure, which is inherent of the respective data assimilation method.

4.1.3   Ensemble Kalman Filters

Ensemble filter methods are a widely used data assimilation techniques. They use the statistical properties of a forecast ensemble to estimate the temporally and spatially varying covariance of the background error (Fujita et al., 2007; Anderson et al., 2009). The ensemble forecasts are treated as a random draw from the probability distribution of the model's state given all previously used observations. Ensemble Kalman Filters (EnKF) are appealing because they provide flow-dependent error statistics and therefore are expected to produce an
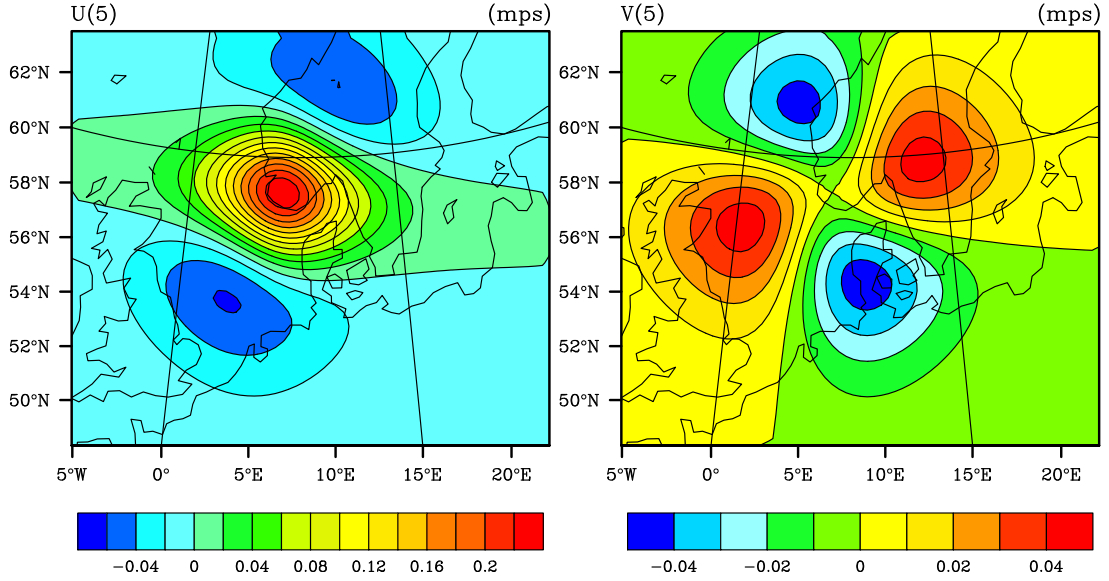
Figure 4.1: Area increments when assimilating a pseudo observation of the horizontal component of the wind (u) of $1\,\mathrm{m\,s^{-1}}$ for u-wind (left) and v-wind (right) on model level 5.
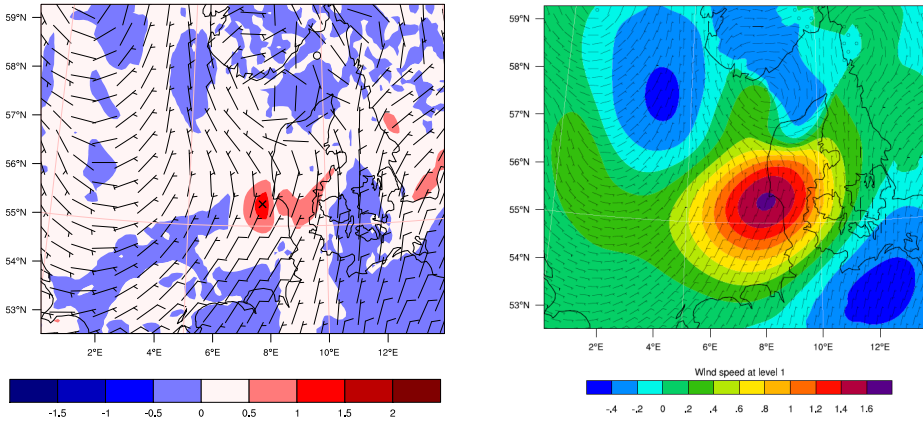
Figure 4.2: Area increments ($\mathrm{m\,s^{-1}}$) when assimilating a single wind turbine observation into FDDA (left) and 3D-VAR (right). The cross in the left figure denotes the location of the wind observation which was also used on the right.

analysis with smaller errors than schemes employing fixed-error statistics (Torn and Hakim, 2008). Using that *flow-dependent* background error covariance can be regarded as the major difference to FDDA and variational approaches. Furthermore, simple implementations require little effort and expert knowledge (Anderson, 2007a). Meanwhile, the EnKF is used in operational settings (Torn and Hakim, 2008). Although it is not trivial to match variational assimilation performance with ensemble methods (Anderson, 2007a), assimilation

experiments with an EnKF system show improvements compared to assimilations with a variational data assimilation system (Torn and Hakim, 2008; Meng and Zhang, 2011). EnKF have been shown to work well in synoptic-scale to convective-scale applications (Fujita et al. (2007) and references therein).

4.2   Data assimilation for wind energy predictions

As will be explained in more detail in Chapter 5, with an increasing number of wind farm deployments, mostly offshore, a new source of data is becoming available especially in data scarce areas: nacelle winds, measured on the nacelle of a wind turbine, and yaw angles, the angle that describes the rotation of the nacelle into the wind. In this thesis the method of WRF FDDA is used with an attempt to improve analyses that are used as initial conditions for short-term wind predictions for wind energy purposes. As a novel approach nacelle winds and yaw angles were assimilated with the aim of improving wind predictions at a downwind location and at the wind farm itself.

Data assimilation in the PBL

Since wind turbines are situated in the boundary layer, the assimilation of wind farm data faces the challenges of data assimilation in the boundary layer; in fact, data assimilation in the PBL is a research topic in itself.

The high degree of spatial and temporal heterogeneity at the surface and in the PBL leads to data representativeness problems that contribute to the difficulty of using these observations (Reen and Stauffer, 2010; Stauffer et al., 1991). It is thus difficult to consider observations near the surface or in the PBL in data assimilation in the context of any dynamic balance, because of the dominance of local forcing (Warner, 2011).

A question addressed in research is the depth in the atmosphere through which the observations should be applied (Reen and Stauffer, 2010; Stauffer et al., 1991; Benjamin et al., 2009; Liu et al., 2008). This is important, because vertical mixing in the model can quickly eliminate near-surface information that is incorporated in the initial conditions, if the atmosphere above is not analysed with vertically constistent structures (Warner, 2011). Hacker and Snyder (2005) and Hacker and Rostkier-Edelstein (2007) demonstrated that surface observations may affect the entire modeled PBL, and demonstrated with an EnKF system that surface observations can substantially improve initial conditions and wind profiles in the PBL.

Meng and Zhang (2011) summarize the problems of limited-area data assimilation, with special focus on Ensemble Kalman Filters, and mention that, because of the error from the difference between the real and the model terrain height and uncertainties in the parameterization of boundary layer and land surface physical processes, surface observations have been challenging in the mesoscale data assimilation field. Moreover, although most statistical data assimilation methods assume that the model forecast (or first guess) is unbiased, that is rarely the case. Model bias error can systematically cause the model to drift away from the truth.

The EnKF has a strong dependence on the accuracy of the model (Anderson, 2007a). In the PBL a significant portion of model error is due to uncertainties in PBL parameteriza-

tions (Hacker and Snyder, 2005). When dealing with assimilation in the PBL, error growth dynamics are often non-linear and non-Gaussian, requiring dedicated treatments compared to synoptic-scale applications (Meng and Zhang, 2011). Since an EnKF system depends critically on the quality of the first guess and the forecast error covariance estimated from a short-term ensemble forecast, the presence of model error is critical. When model error is not included in the ensemble-based calculation of forecast error covariance, the resulting spread may be insufficient. Ways of dealing with imperfect models in the PBL include covariance inflation procedures, bias correction algorithms, stochastic approaches, or the use of multi-physics ensembles. One other reason why EnKF have issues in the PBL is due to sampling error. Localization is a means to ameliorate sampling error when small ensembles are used. The EnKF methods use linear regressions. Sampling error in linear regression can be the dominant source of error in the whole filtering procedure (Anderson, 2007b). Localization, especially in the vertical, is challenging in the PBL (Mitchell et al., 2002), since the area of influence of an observation should change with atmospheric stability. In the presence of severe model error, the radius of influence should be smaller.

The assimilation of wind farm data, specifically, also touches upon the problem that datasets with a high spatial density violate the assumption of spatially independent observation errors in data assimilation (e.g., Daley (1991), Ochotta et al. (2005)). Wind farm data are spatially dense data: compared to a mesoscale model grid with grid spacings between 2 and 30 km, many observations exist within one grid cell. Therefore a method is needed to reduce the amount of data and produce a representative data set, i.e., to thin the data.

If only data near the surface are assimilated, the future model integrations could, through geostrophic adjustment, "wipe out" the changes that were done to the analysis via the assimilation. This suggests that upper air observations are needed in assimilation systems to keep the impact of near surface observations in the system (Barwell and Lorenc, 1985; Stauffer et al., 1991). Additionally, it can be expected that in locally forced weather situations, the effect of the assimilation will stay in the system for longer as compared to synoptically forced situations. To give an example to the latter, a cold front passing through just after the time of the assimilation will likely "wipe out" the data assimilation effect.

4.3   Conclusions

This chapter reviewed the basic concept of data assimilation, the methods of FDDA, variational data assimilation and EnKF. The idea of assimilating data in the PBL to improve wind energy forecasts is discussed. Even though the different data assimilation systems have each their advantages and disadvantages and some might be more suited in certain situations than others, the challenges of introducing a new data set to be assimilated are shared by all of them. The methods and issues relating to the assimilation of wind farm data explained on the example of FDDA in the following chapter can thus be applied to any data assimilation system. With more and more wind farms being built and wind farm observations becoming increasingly available, adding these to the already existing data used operationally to improve the initial conditions of a NWP model, will be promising for both wind energy predictions and weather predictions in general. Since wind farm operators are reluctant to share their data, a lot of lobbying will be needed to allow wind farm data to be widely used. Additionally, further research than presented in this thesis will be necessary for wind farm

observations to be considered potential candidates in contributing to a richer data set than is currently used in NWP. This thesis is a first step towards this endeavour.

CHAPTER 5

Wind Farm Data

This chapter serves as an introduction to the area of study and data used for the data assimilation experiments in the next chapter. Reference is given to the Horns Rev I wind farm and its turbines, Vestas V80, but an effort is made to generalize the issues of wind farm data, so that the ideas presented here can be exploited for other applications, purposes and wind farms. The explanation of the data is mainly from the point of view of their use in data assimilation.

5.1 Area of study - The wind farm Horns Rev I

The area of research for the assimilation experiments is the wind farm Horns Rev I (Figure 5.1). The wind farm consists of 80 turbines, which are distributed in 10 rows with 8 turbines each on an area of 19.7 km^2, approximately a 5 x 3.8 km square. The distance between the turbines is 560 m, corresponding to seven rotor diameters. The distance from the coast (to Blåvands Huk) is 13.8 km. Halfway between the coast and the wind farm, 6 km from the farm, a measurement tower (referred to as M7) is installed. During westerly winds, this measurement tower is in the wake of the wind farm. One measurement mast is situated 2 km to the northwest of the farm (referred to as M2), and in the wake of the wind farm during south easterly winds. Data used from M7 are wind speeds at 70 m and wind direction at 68 m. From M2 I use wind speed measurements at 62 m and wind direction measurements at 43 m. Wind anemometers are top-mounted Risø cup anemometers, and therefore not subject to tower and boom distortion. The wind direction measurements from M2 have a known offset. I use these wind direction measurements in one experiment, further pressure measurements from M2 at 55 m and temperature measurements at 13 and 55 m, to calculate the pressure at 70 m with the hydrostatic assumption.

Measurement masts are primarily set up in a wind farm to evaluate the wind potential at a site before the wind farm is built. This was also the case for M2. Reliability of these data can therefore degrade if they are not as well maintained after their initial purpose. The data
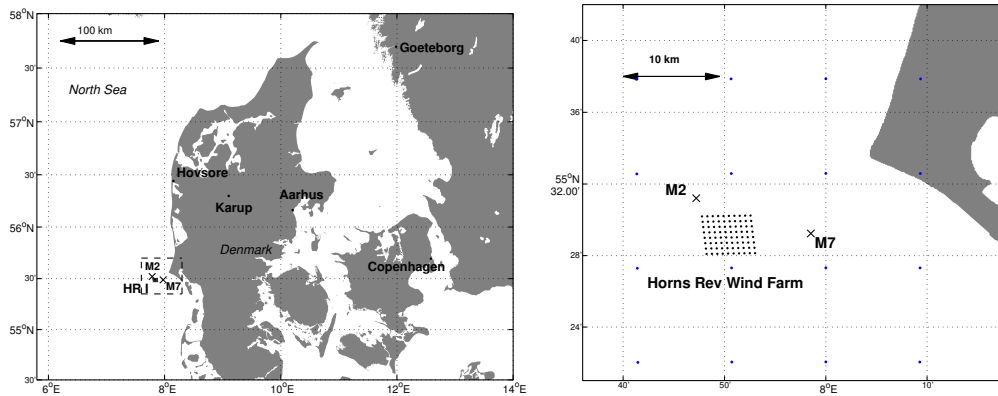
Figure 5.1: Map of Denmark (left); The filled square denotes the location of the wind farm Horns Rev I, the crosses the location of the measurement masts M2 and M7 respectively. The dashed square indicates the area of the map to the right, where the dots represent the model mass grid points of domain 2 (used in Chapter 6).

from M2 show good reliablity most of the time. M7 was erected together with another mast situated between M7 and the wind farm mainly to observe the wind farm wake when the wind comes from westerly directions (Hasager et al., 2007).

The hub height of the turbines is 70 m MSL, which corresponds to the height used in the data assimilation experiments. This height is not constant but changes with sea level. These changes occur due to swell and tides, and are in the order of 1.31 m below DNN (Danish Normal Zero) to 2.8 m above DNN (Sommer, 2002) at Horns Rev.

5.2 Wind speeds

Each wind turbine in a wind farm is equipped with anemometers that are located on the nacelle behind the rotor. Figure 5.2 illustrates the location of the nacelle and nacelle anemometers of an example turbine. Depending on the turbine manufacturer, the distance of these anemometers from the rotor varies. In the case of Horns Rev I, ultrasonic anemometers are situated at the end of the 10.4 m long nacelle, measuring linear components of the wind vector u and v. These wind measurements, i.e., nacelle winds, are usually used by wind farm operators for the turbine control, to determine cut-in/cut-off wind speeds. If nacelle winds are used for other purposes, like in this case for data assimilation, knowledge about the issues associated with these measurements is important. Since the anemometers are situated behind the rotor, they do not represent the free wind conditions from that region but are disturbed by the turbine and rotor blades. This disturbance of the measured wind depends on the design/shape of the wind turbine and nacelle, the pitch/stall situation, whether the turbine is operating or not, the height of the anemometer and its position on the nacelle (Zahle and Sørensen, 2011; Smaïli and Masson, 2004). For data assimilation purposes the information from the free wind conditions at a site are required, and deviation from representative wind speeds lead to errors. Since the wind disturbances are very complex, a general valid correction of these data is an ongoing field of research.
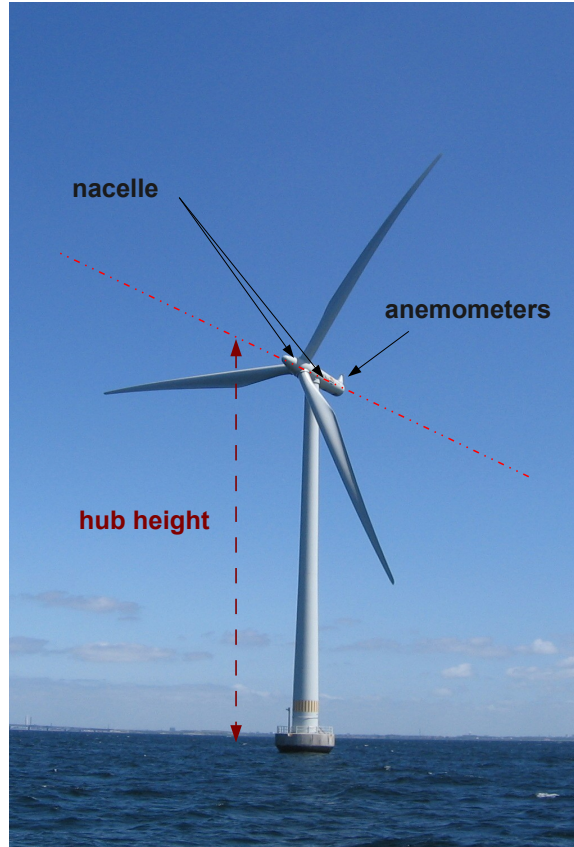
Figure 5.2: Photo of a wind turbine to demonstrate the nacelle, the location of nacelle anemometers and the hub height. This turbine is part of the Swedish offshore wind farm Lillegrund.

However, nacelle anemometers are calibrated with a so-called nacelle transfer function to describe the undisturbed wind field (Frandsen et al., 2009), therefore part of the above mentioned issues are partially already accounted for. This function is determined empirically by comparing wind measurements from a tower that is approx. 250 m upwind with the output of the nacelle anemometers. The function exists for both wind speed and direction but is still limited and does not take turbulence effects, vertical inflow or upflow conditions into account. Furthermore, these functions are based on a reference turbine, and not calculated for each specific turbine. Therefore, they do not account for local conditions or park effects (Zahle and Sørensen, 2011; Antoniou and Pedersen, 1997). However, it can serve as a good approximation during *normal* conditions, e.g. for reasonable amounts of turbulence or no yawing errors (i.e., the difference of wind direction of the free wind to the correct alignment of the turbine in the wind) (Antoniou and Pedersen, 1997).

Most turbines are equipped with two anemometers, one on the right and one on the left on the nacelle. Sometimes, when data transfer problems or data overload occur, the received signal can switch from one anemometer to the other one. Between the two anemometers can thus be a small difference, adding a further factor of complexity to the issue (Rolf-Erik Keck, Vestas, personal communication, Dec. 2010).

5.3 Wind direction

In data assimilation, the horizontal wind components, u and v, are assimilated. Therefore wind direction measurements are required. Wind turbines, in fact the whole nacelle with the rotor, is turning itself into the wind based on wind direction measurements from behind the rotor. This rotation, the yaw angle, is recorded and available as a proxy for wind direction. Depending on the turbine type, turbines yaw quickly: Common MW-turbines yaw with a yaw-rate of about $0.5\,^\circ\,\mathrm{sec}^{-1}$. They use averaging algorithms, and typically yaw when they have reached an averaged yaw-error value of about $+5\,^\circ$ or $-5\,^\circ$. Then they yaw until they have reached the center value again. As with nacelle winds, yaw angles also have associated problems.

First of all, the wind direction which is measured with anemometers situated on the nacelle to determine the yaw error is error prone (Zahle and Sørensen, 2011). Yaw errors appear to be on average up to 10 degrees (Pedersen et al., 2010) and are a result of the difference between the free wind direction and the measured direction of the disturbed wind behind the rotor.

Secondly, below the cut-in wind speed of 4 $\mathrm{m\,s}^{-1}$ and above the cut-off wind speed of 25 $\mathrm{m\,s}^{-1}$, when the turbine is not operating and thus not yawing, the yaw signal is apparently wrong. Upon restart, with the turbine then being in a random position, it takes time for the sensor to be oriented towards north again, resulting in a messy data set. Further complicating is the fact, that the turbines' cables have to be disentangled after too much rotation of the nacelle, and the yaw has to be oriented towards north again.

Thirdly, each wind turbine has a different offset in yaw angle. Within a recent EU project that looked towards the design of very large wind turbines (8-10 MW) both onshore and offshore (the project UpWind), the offset of the yaw angle for Horns Rev's turbine 7 was corrected by determining the power deficit between two turbines as a function of direction (Kurt S. Hansen, personal communication, Dec. 2010).

Despite these potential problems, comparisons of yaw angles with wind direction measurements from masts show quite good agreement during the experiments (indicated by Figure 5.3 for one case), so yaw angles from turbine 7 are used as a proxy for wind direction in the experiments in the following chapters. The timeseries of the yaw angle are smoother than the ones from mast M7 downwind, which might even be an advantage in the assimilation process, because they might fit better to the hourly model values and thus not be eliminated by the data misfit or quality control procedures. The lag shown with respect to M7 is probably due to the distance of 6 km, indicating that the turbine reacted to wind direction changes *before* they hit M7 downstream.

5.4 Wake effects

As well as the disturbance of the wind due to the turbine's rotor, wind measurements on turbines in a wind farm are affected by wakes. Losses in turbine power output due to wakes can be in the order of 5-15% of a whole wind farm (Barthelmie et al. (2010), Barthelmie and Jensen (2010)) and the wind speed decrease leading to these power losses can be up

Figure 5.3: Yaw angles and wind directions from measurement mast M7 for 29 May 2005 as an example for a typical time series of the data.

to 20% for certain turbines depending on their position in the wind farm with respect to the flow direction, which can result in a few $\mathrm{m\,s^{-1}}$. Barthelmie et al. (2010) state, however, that wake modeling of large wind farms is still subject to an unacceptably high degree of uncertainty. The magnitude and extent of wakes depend on atmospheric stability and wind speed (Christiansen and Hasager (2005b), Barthelmie et al. (2010), Hansen et al. (2012)). Furthermore, wind farm layouts and farm sizes have an impact as well.

Figure 5.4 shows timeseries of nacelle winds for the first five upwind turbine rows and the 70 m measurements of mast M7 on 29 May 2005. The wind hit the farm from the west; The wind speeds from the mast are lower than the ones from the turbines in row 1, but the further to the east the row is situated, smaller differences between the mast measurements and the turbine measurements can be seen. The further east the wind farm rows are situated, the more wake effects they experience. This indicates a slight wake effect on mast M7 since its measurements correspond best to the wake contaminated nacelle winds from row 5. The 70 m measurements from M7 downwind show a lag around 14:30 UTC, which can be explained by the distance of 6 km between the farm and the mast. The nacelle winds from each row are within an interval of approximately 1 $\mathrm{m\,s^{-1}}$.

Correction of nacelle winds using WAsP

Several approaches exist to get rid of the above mentioned effects and clean the data to represent free wind speeds or wake free winds. One way of "cleaning" the dataset from wake effects can be done by Risø's microscale model WAsP Engineering (Beyer et al., 1994). It estimates wake velocity deficits and is based on a linear wake expansion of the wake downstream (originally developed by Jensen (1983)) and momentum balance in uniform flow.

The relative deficit at a rotor hit by a single wake is estimated by

$$\Delta u_1/u_1 = (1 - \sqrt{1 - C_T(u_0)}\frac{A_{0,rotor}}{A_{1,wake}}\frac{A_{1,overlap}}{A_{1,rotor}} \tag{5.1}$$

Here, u_0 and u_1 are the wind speeds at hub-height at the upwind and downwind turbine sites, respectively. $A_{0,rotor}$ and $A_{1,rotor}$ are the rotor swept areas of the upwind and downwind turbines, $A_{1,wake}$ the area of the expanded wake at the downwind position set with a wake decay factor, and $A_{1,overlap}$ the overlap area of the expanded wake and the exposed rotor. Δu_1 is the resulting velocity deficit at the downwind turbine site.

Since the wake is assumed to expand linearly behind the rotor, wake meandering (Larsen et al., 2007; Espana et al., 2009) or small scale local effects can not be corrected for with WAsP Engineering. The resulting correction factors for each turbine depends on the incoming wind direction and the wind speed.

The absolute differences between WAsP corrected nacelle wind speeds versus non-corrected ones for the assimilation experiments carried out in this thesis in Chapter 6 did not exceed 1.6 $\mathrm{m\,s}^{-1}$ and a qualitative assessment of results assimilating WAsP corrected nacelle wind speeds versus non-corrected ones revealed only minor differences (not shown). This method was thus not exploited further and nacelle winds were therefore not corrected for the simulations in Chapter 6.

5.5 Wind farm data in data assimilation

In every discipline, one has to know the data set one is using. Depending on the usage of wind farm data, they have to be processed to be meaningful. Here I mention what has to be considered when these data are used in data assimilation.

As already mentioned in Chapter 4, a method is needed to reduce the amount of data and produce a representative data set. Since in my results in Chapter 6 I focus on quantifying the hypothetical impact of the wind farm data assimilation downwind at M7, as well as on the wind farm itself, error metrics were calculated using the wind measurements from M7 as a reference. Data and their timeseries were analysed and medians of different turbine groupings with respect to M7 calculated. I preferred the median over the mean to avoid the negative effect of outliers. The groupings serve as a thinning technique (section 4.2). Performing different experiments by using the different groupings is also a way to deal with the inherent uncertainties of wind farm data.

In order to keep the impact of wind farm data in the system (Chapter 4) and also because certain observations are considered standard observations, wind farm data were assimilated together with upper air observations of the Meteorological Assimilation Data Ingest System (MADIS, http://madis.noaa.gov), which include data from the Aircraft Communication Addressing and Reporting System (ACARS), and observations from radiosondes, to analyse the effect of wind farm data in data assimilation.
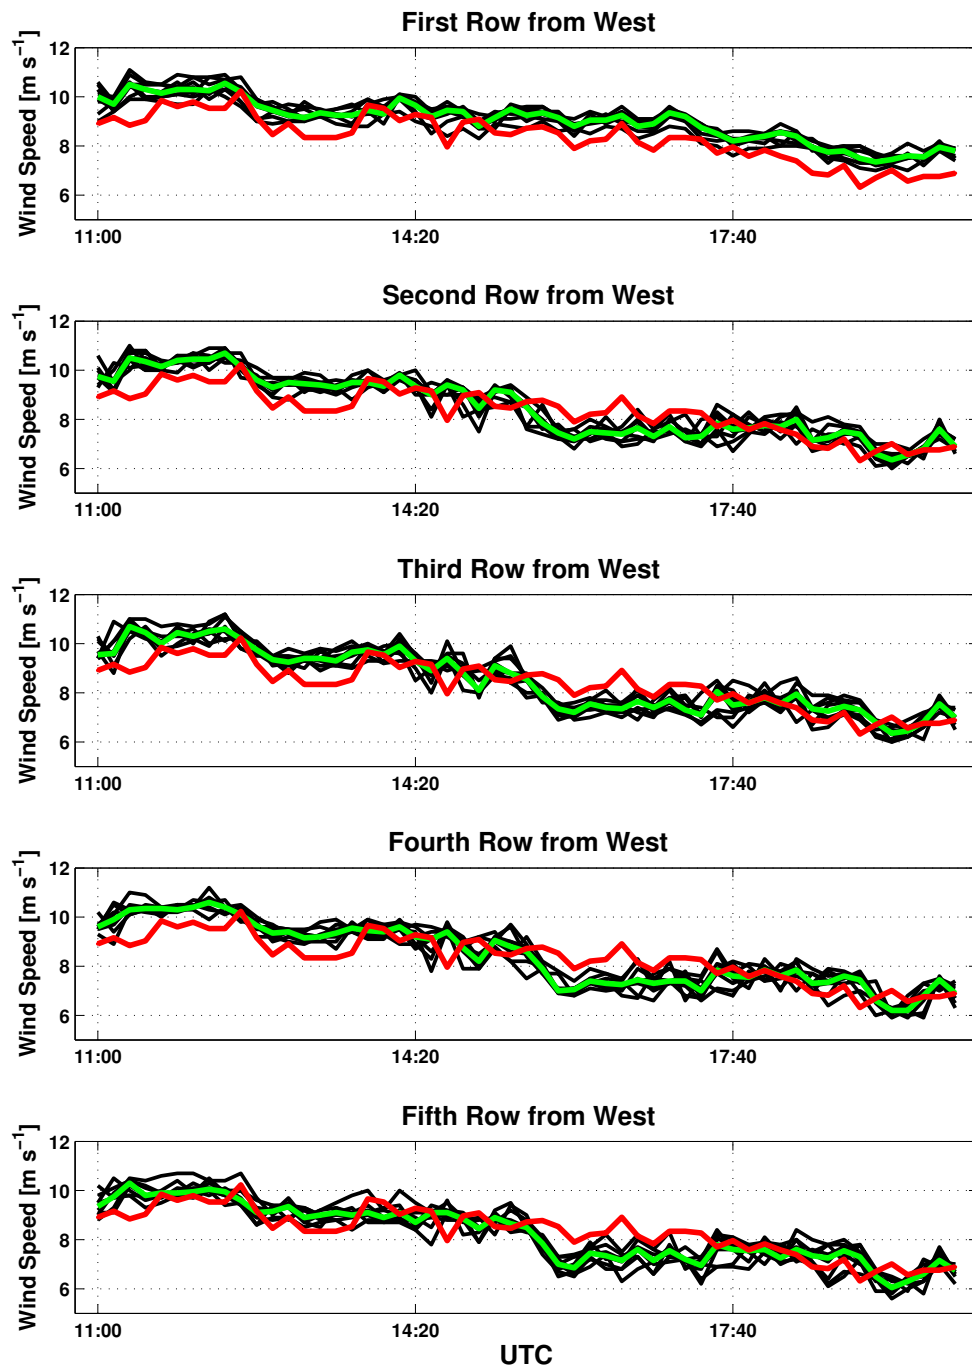
Figure 5.4: Timeseries of nacelle winds (individual turbines in black, median in green) and wind speed measurements from measurement mast M7 at 70 m (red) for each turbine from the westernmost row of the wind farm (upper panel), the second row from the west (second panel), the third row (third panel), the fourth row (fourth panel) and the fifth row from the west (fifth panel). The wind on this day (29 May 2005) came directly from the west.

In this thesis I use wind farm data as described above, but other methods exist to make use of the data. For the sake of completeness they are mentioned here:

Instead of using yaw angles and pressure measurements from a measurement mast, model values for wind direction and pressure could be used. Pressure values are needed in the assimilation to assign the correct height to the measurements.

Wind direction is also available from measurements on the nacelle, in the case of Horns Rev I from sonic anemometers. However, whenever the turbines are reset, the reference angle determining north becomes incorrect. It takes a while until it goes through a reference direction before the sensor is oriented towards north again. This results in poor data quality. Furthermore, this angle is also modified by the flow through the rotor.

Assimilating power data is a possibility, but power data are very costly and wind farm operators are usually very reluctant to distribute them, making use of wind farm power data in an operational context a difficult endeavor. By using nacelle winds and yaw angles I made an effort to suggest a fairly general technique. In some jurisdictions, the SCADA data has to be delivered to the Transmission System Operator, who could use them for wind energy forecasting. For the assimilation of power data an observation operator would need to be developed. This operator would not only depend on physical constraints, but also on economics. The availability of power data is linked to the shut off of wind turbines, which does not only happen due to environmental conditions. Market driven decisions determine the operation of wind farms and it is nearly impossible to determine the reason for shut off on a regular basis. Moreover, these decisions depend on the country or wind farm operator, so no general valid method will be available. Nacelle winds are being measured on the turbines in a farm all the time, including during curtailments.

In addition, it is extremely difficult to relate power to wind. The conversion from power to wind will only be reliable where the slope of the power curve is sufficient, i.e. between approximately 5 and 13 $\mathrm{m\,s^{-1}}$. Furthermore, pressure is needed to calculate the air density. Pressure sensors are usually not part of the standard intrumentation, so this needs to be available from an external source. Using nacelle winds and yaw angles as in this study is a fairly feasible and promising approach for the future, once wind farm data will be used more frequently in data assimilation systems.

5.6 Conclusions

This chapter shed light on the area of study in the data assimilation experiments in the next chapters. The wind farm data used are explained and their issues discussed. It is very clear, why wind farm data are a data source that is questioned by some when it comes to data assimilation. Understanding the data set and accounting for the issues is a first step to alleviate its issues and to make the best use of the data. It is also clear that not every wind farm is as well maintained and documented as Horns Rev I, which might make it challenging to adapt my approaches to other cases. While most offshore wind farm operators are professional enterprises, which ensures a certain data quality, onshore data from onshore wind farms can lack that quality. Assimilation results are presented in the next chapter, as a proof of concept, that this data set can be successfully used.

CHAPTER 6

Wind Farm Data Assimilation in FDDA

This chapter uses the information from Chapter 5 and presents the impact of the assimilation of wind farm data into the Four-Dimensional Data Assimilation (FDDA) system of the WRF model (Skamarock et al., 2008). Wind farm observations, which up to now are primarily used by wind farm operators for control purposes, constitute a novel and unique set of measurements in the PBL.

Not many contributions can be found in the literature about the use of nacelle winds in mesoscale forecasting, none about yaw angles, to the best of my knowledge. Cutler et al. (2011) used nacelle wind speeds for power curve modelling and suggest the use of nacelle winds averaged over a wind farm in wind power forecasting. Kankiewicz et al. (2010) presented a poster on the use of nacelle anemometers for forecasting and operational assessment. Assimilating wind speeds measured on the nacelle of a wind turbine and the yaw angle, which serves as a proxy for wind direction, into mesoscale models is a novel topic. While data assimilation experiments have been done with various other measurements to improve the forecasts for wind farms (Liu et al., 2011; Delle Monache et al., 2010; Zupanski et al., 2010), only few conference contributions exist from the National Center of Atmospheric Research (NCAR) about the assimilation of tower measurements in the vicinity of onshore wind farms and nacelle winds using NCAR's non-community Real-Time Four Dimensional Data Assimilation (RT–FDDA) model (Cheng et al., 2011; Liu et al., 2010). To the best of my knowledge, yaw angle observations have not been used for data assimilation purposes to this date.

Since the measurements on the individual wind turbines are not independent from each other and dealing with each turbine's measurements separately makes the data handling challenging, thinning the data is desireable. The focus of section 6.1, the content of which is part of the paper Draxl et al. (2011a), is thus on exploring different thinning strategies and their impact on four-dimensional data assimilation concerning wind predictions at hub height. The experiments analysed here are from hourly analyses, i.e., the results are from forecasts that are initialized one hour apart from each other. This is contrary to most current operational weather forecast systems, that initialize a new forecast every 6 hours. However, on the small time scales relevant to predicting wind speeds for wind energy applications, the

local scale development of systems may need to be updated more frequently. Since FDDA is a relatively computationally efficient technique, hourly updates are feasible. As the basics for wind farm data assimilation and their associated issues are similar in any data assimilation system, section 6.1 serves as a guideline for the assimilation of wind farm data in any data assimilation system.

6.1 Thinning strategies for the assimilation of wind farm observations

Today many different kinds of observations are used in operational numerical weather prediction (NWP) models with the aim of more accurate forecasts by improving the analysis and subsequent forecast through data assimilation. In this section I will explore the potential benefits of the assimilation of wind farm observations in addition to the already available upper air data sets used for data assimilation. The area of study is the Danish offshore wind farm Horns Rev I. The purpose of this study is to explore different thinning strategies and their impact on analyses and short-term predictions.

6.1.1 Data

The wind farm data used for the assimilation experiments, nacelle winds and yaw angles, are described in Chapter 5. The wind farm Horns Rev I is located off the west coast of Denmark in the North Sea (Figure 5.1). In the assimilation experiments a constant height of 70 m MSL is assumed for the observations.

Halfway between the coast and the wind farm (i.e., 6 km to the east of the farm) a measurement tower (referred to as M7) is installed. During westerly winds, this measurement tower is in the wake of the wind farm. Another measurement mast is situated 2 km to the northwest of the farm (referred to as M2), and in the wake of the wind farm with southeasterly winds. The nacelle winds, yaw angles and tower data are available as 10 min averages.

Data used from M7 are wind speeds at 70 m and wind direction at 68 m. Wind speed measurements at 62 m and wind direction measurements at 43 m are used from M2. I further used wind direction from M2 in one experiment as a comparison to using only yaw angles (section 6.1.5). Pressure measurements at 55 m and temperature measurements at 13 and 55 m are taken to calculate the pressure at 70 m with the hydrostatic assumption, which is needed for the assimilation.

6.1.2 Data assimilation system

The Four Dimensional Data Assimilation system of the WRF model is based on Newtonian relaxation (i.e., nudging). I used the approach of nudging directly toward individual observations. This approach uses only those observations that fall within a predetermined time

Figure 6.1: Domain configuration and terrain elevation of the WRF model setup. The squares indicate the boundaries of nested domain and the cross the location of the wind farm Horns Rev I.

Figure 6.2: Schematic for the setup of the data assimilation experiments.

window centered about the current model time step. In this set up the half time window for domain 1 was 0.8 hours, for domain 2 it was 0.6 hours.

The model setup for the experiments consisted of an outer and nested domain with horizontal grid increments of 30 and 10 km, respectively (Figure 6.1). The model was initialized and forced at the boundaries by $1° \times 1°$ U.S. National Center for Environmental Prediction (NCEP) Global Forecast System 6-hour forecasts. The sea surface temperature fields were also obtained from NCEP analysis at a horizontal resolution of $0.5° \times 0.5°$. 1-way nesting and 37 vertical levels were used, with 7 levels within the lowest 500 m. The lowest levels important for wind energy applications were at approximately 16, 59, 117, and 183 m AGL. The model physics options include: Yonsei University (YSU) boundary layer scheme, Lin et al. microphysics scheme, MM5 similarity surface layer, Noah land surface model, and the Kain-Fritsch cumulus parameterization. The model dynamics options include: 2nd-order diffusion with a horizontal Smagorinsky first order closure, 6th-order numerical diffusion with prohibit up-gradient diffusion for the inner domain, and positive-definite advection of moisture and scalars. The radius of influence in domain 1 was 200 km and in domain 2 it was 90 km.

Figure 6.3: An example for the distribution of observation locations on 20 July 2005 at 12 UTC. Radiosonde observations (red cross) and ACARS (green dot) were used in the assimilation. The black cross denotes the location of Horns Rev.

Observations were assimilated hourly during a pre-forecast period of 24 hours to provide spun-up initial conditions for each of the test cases described in section 6.1.3, and for the time of the test cases thereafter to produce initial conditions (Figure 6.2). The control run where no data assimilation was performed (CTL) was spun up without nudging. After the spin up period, free 24-hour forecasts were initialized from the resulting analyses every hour. For the selected cases this resulted in 44 24-hour forecasts. The assimilated data sets included wind farm data and upper air observations of the Meteorological Assimilation Data Ingest System (MADIS), which include data from the Aircraft Communication Addressing and Reporting System (ACARS), and observations from radiosondes; a snapshot of their typical distribution is shown in Figure 6.3.

6.1.2.1   Thinning strategies

Wind farm data are spatially dense and compared to a mesoscale model grid with grid spacings between 2 and 30 km many observations are in the subgrid space (Figure 5.1, right panel). Thus, the measurements do not provide independent information and dealing with each turbine's measurements separately makes data handling challenging and increases the computational costs (Lazarus et al., 2010). Additionally, in data assimilation systems like the Ensemble Kalman Filter or variational methods, datasets with a high spatial density violate the assumption of spatially independent observation errors (e.g., Daley (1991) or Ochotta et al. (2005)). Therefore a procedure is needed to reduce the amount of data and

produce a representative data set. Thinning is also a way to deal with the inherent issues associated with wind farm data (section 5.2 and 5.3).

The thinning strategies consist of different grouping strategies of nacelle winds. In this study, I calculated median nacelle winds of the first upwind turbine row, the second upwind row, the third upwind row, the fourth upwind row, the fifth upwind row, all the turbines, the upwind half of the wind farm and the first two upwind turbine rows combined. I chose the median over the mean to avoid the negative effect of outliers. Nacelle winds were quality controlled, but not corrected for their inherent issues, as mentioned in section 5.4. For the assimilation experiments, wind directions come from the yaw angle of turbine 7 (situated in the first upwind row), because data quality of that turbine was assured (section 5.3), and additionally from measurement mast M2 at 60 m for the group with medians of the first two upwind rows.

The metrics used to evaluate the best thinning strategies were the root-mean-sqared error (RMSE), bias, centered root-mean-sqared error (CRMSE) and rank correlation (section 6.1.4).

6.1.3   Test cases

Suitable test cases were selected to study the impact of the different thinning strategies and of the assimilation of wind farm data following these strategies. The impact of the assimilation at the location of the observations, i.e., the wind farm, as well as at downwind locations is of particular interest. The latter is important for the wind energy industry to understand if data collected at a wind farm could be useful for short-term predictions at a nearby downwind farm. This is the case at Horns Rev I (Figure 5.1), which is situated downwind of Horns Rev II during westerly winds. Wind farm data from Horns Rev II were not available at the time this research was performed. Given the position of M7 with respect to Horns Rev I cases with westerly winds (i.e., winds between $260\,°$ and $280\,°$) were selected. Restricting the cases to westerly winds would also allow us to define unchanging upwind rows.

However, during westerly winds, M7 could be affected by wakes of the wind farm. The selected periods were thus further restricted to periods where wind speeds at M7 were within $\pm\,5\%$ of mast M2, the mast to the northwest of the farm. This was to ensure impartial verification with minor wake effects at M7, which is in a wake during westerly winds. Results with Risø's wake model WAsP Engineering (http://wasp.dk, last access: Feb. 2012) confirmed minor wake effects on M7 below $1.6\ \mathrm{m\,s^{-1}}$ during the test cases because of the distance to the wind farm (not shown).

Due to the restrictions to westerly winds and to periods with minor wake effects, only 5 cases with varying length in May, July, August and October 2005 were found with simultaneous measurements on the turbines and the measurement masts out of ~ 10 years of recorded data. Given that I analyzed 24-hour forecasts the restrictions will last only during the time of the analyses and at least for the first three forecast hours. For later lead times wind conditions may change.

Table 6.1: Date and weather conditions of the test cases. The dates denote the analyses that were used to initialize the forecasts.

Test case	Weather conditions at Horns Rev
29 May 11 UTC - 29 May 17 UTC	A low pressure system over the British Isles and southern Scandinavia led to southern winds in 500 hPa and westerly winds near the surface. During the course of the day high pressure took over from the south, leaving Horns Rev all day in a westerly flow
19 July 13 UTC - 20 July 05 UTC	Low pressure dominated Scandinavia and the British Isles, with its core moving from the north of the British Isles westward, leaving Denmark to the south of it in a westerly flow. A trough axis passed over western Denmark between 6 and 12 UTC, just before the test case. There was precipitation associated during this case
26 Aug. 00 UTC - 26 Aug. 04 UTC	Low pressure was dominating over Northwest Europe, leading to westerly flow with convective conditions at Horns Rev
26 Aug. 13 UTC - 26 Aug. 18 UTC	Low pressure was dominating over Northwest Europe, leading to westerly flow with convective conditions and showers at Horns Rev
25 Oct. 16 UTC - 26 Oct. 00 UTC	Low pressure was dominating over northern Europe, with its center located over southern Norway at 18 UTC. A cold front was associated with the low, which had passed Horns Rev around 12 UTC on October 25th. Denmark was in a westerly flow and precipitation was associated with this case

Denmark is situated around 56 degrees north with prevailing westerly winds. The selected test cases with winds from the west are thus representative. During the experiments the conditions at the wind farm Horns Rev were associated with a low pressure system over northern Europe/Scandinavia and/or a cold front passage (Figure 6.4, Table 6.1). Unstable atmospheric stability conditions prevailed during the test cases. The YSU boundary layer scheme was thus the PBL scheme of choice, since it was found to perform best during unstable cases in the boundary layer comparison study in section 3.3.

6.1.4 Verification metrics

For wind energy forecasting both the prediction of the timing and amplitude of an event are important. Capturing the amplitude is crucial, because the power is a function of the cube of the wind speed. Timing is important during ramp events, i.e., large and rapid changes in wind speed, which can only be predicted well if the pattern variablities are forecast correctly. In this evaluation the focus is on

- Root-mean-squared-error (RMSE): This is the metric that is best known and commonly used. It can be split into the bias and centered root-mean-squared-error (CRMSE), which are the systematic and random components of the RMSE, respectively (Taylor, 2001).

$$RMSE^2 = \frac{1}{N_p}\sum_{i=1}^{N_p}(F_i - O_i)^2 = CRMSE^2 + BIAS^2 \qquad (6.1)$$

N_p is the number of available forecast(F)-observation(O) pairs.

Figure 6.4: Maps of sea level pressure during the test cases; for 29 May 2005 at 15 UTC (upper left), 20 July 2005 00 UTC (upper right), 26 Aug. 2005 at 12 UTC (lower left) and 25 Oct. 2005 at 18 UTC (lower right). Overlaid are wind barbs on the second model level (~59 m), half ticks representing 5 m s^{-1}, full ticks 10 m s^{-1}. The cross denotes the location of Horns Rev I.

- BIAS: The bias shows the differences in the means of the observations ($\overline{O}$) and the forecasts ($\overline{F}$) and is a measure for the systematic component of the forecast error, which may come from the model misrepresentation of topography/coastlines, offset parameter values, biased initial conditions, etc. (Delle Monache et al., 2011):

$$BIAS = \overline{F} - \overline{O} \tag{6.2}$$

$\overline{F}$ and $\overline{O}$ are the forecast and observation averages over N_p values.

- CRMSE: The CRMSE is considered the random component of the error. It gives an indication of the intrinsic predictive skill of the forecast that can be limited by the coarse or nonexistent representation of specific physical processes (Delle Monache et al., 2011), which is often harder to get rid of in a forecasting system than the systematic bias. The CRMSE penalizes phase errors.

$$CRMSE = \sqrt{\frac{1}{N_p}\sum_{i=1}^{N_p}[(F_i - \overline{F}) - (O_i - \overline{O})]^2} \tag{6.3}$$

- Spearman Rank Correlation (Wilks, 2006): The rank correlation is a nonparametric (i.e., distribution free) statistic that reflects the strength of the monotone relationships between two variables. As such it allows for a nonlinear relationship between the predictions and observations, which is appropriate when the quantity of interest, wind speed in this study, exhibits a non-Gaussian distribution, and quantifies pattern similarity (Delle Monache et al., 2011). High correlation values indicate better pattern similarity. This is especially important in wind energy forecasting where a correct prediction of the onset of a ramp event is extremely valuable. Data are transformed into ranks, i.e., they consist of all the integers from 1 through the sample size N. In the following equation, D_i is the difference in ranks between the i^{th} pair of data values.

$$r_{rank} = 1 - \frac{6 \sum_{i=1}^{N} D_i^2}{N(N^2 - 1)} \qquad (6.4)$$

6.1.5 Results

6.1.5.1 Impact of thinning strategies

As described in section 6.1.2.1, thinning techniques are necessary for wind farm data to be assimilated into a NWP model. In this study they consist in different grouping strategies of nacelle winds. This study focuses on quantifying the data assimilation impact on wind speed forecasts downwind at M7 and at the wind farm itself. I thus calculated error metrics (bias, RMSE, and rank correlation) for different groups of turbines using the wind measurements from M7 as a reference. This serves as an estimation of which group (i.e., thinning strategy) would have the potential to be successful in the assimilation experiments. Thinning is also a way to deal with the inherent issues associated with wind farm data (sections 5.2 and 5.3).

The groups were: median winds of the first upwind row, the second upwind row, the third upwind row, the fourth upwind row, the fifth upwind row, all the turbines, the upwind half of the wind farm and rows one and two together. A schematic of the groups is shown in Figure 6.5. The median was preferred over the mean to avoid the negative effect of outliers.

The groups mentioned above were ranked by bias, RMSE and rank correlation, and these rankings were summed to select a subset including the best thinning strategies. The turbine aggregations that exhibited the lowest RMSE, rank correlation, and bias were used in the assimilation experiments. The group with the lowest combined error metric is the median of row 1 and 2 together, followed by the median of row 1 and the median of all the turbines (Tables 6.2 and 6.3).

The medians of wind speed of row 1 and rows 1 and 2 together are usually higher than the ones for the other groups due to wake effects (Figure 6.6). The median of all the turbines is lower, because it includes the lower wind speeds from the wake affected turbines downwind. Generally, the difference between the groups is within $1 - 2 \, \mathrm{m\,s^{-1}}$. The bias of all the groups is mostly negative, that means that the wind speeds of M7 are higher than the ones from the groups; the correlation is highest for turbine 1 most of the time - which indicates that the wake effect on M7 was very small. Wake meandering will have an impact on these results as well, but I have not attempted to quantify it. For the assimilation experiments, wind

Figure 6.5: Schematic for the different grouping strategies of nacelle winds. Black dots represent the turbines, underlaid is the number of turbines used to calculate the median.
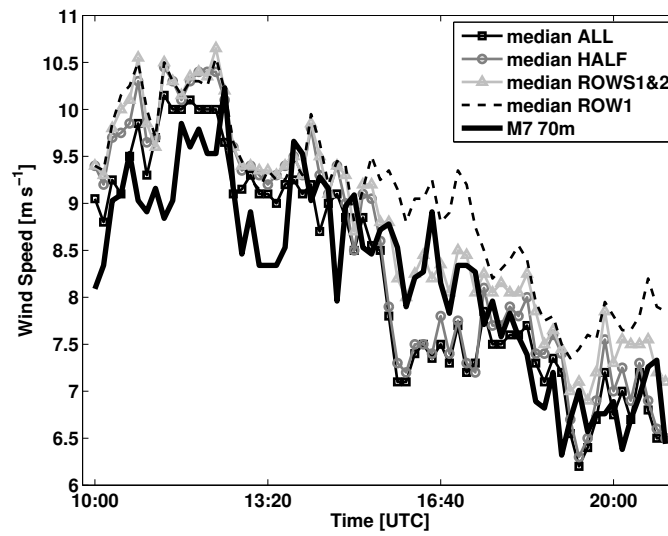
Figure 6.6: Timeseries of median wind speeds of the groups ALL, HALF, ROW1 and ROWS1&2 (Table 6.4) as well as of the 70 m measurements at M7 for the case from 29 May 2005. The median of ROW1 is higher than the others, which is representative for all the cases.

directions come from the yaw angle of turbine 7 and additionally from measurement mast M2 at 60 m for the group with medians of rows 1 and 2 together.

Table 6.2: Root mean square error (RMSE), Pearson correlation and bias with regard to wind speed aggregations with regard to wind speeds at M7 at 60 m, and their ranking for the median winds of the first upwind turbine row, the second row, the third row, the fourth row, the fifth row, all the turbines, half of the wind farm, and rows one and two together for all the test cases separately.

| Aggregation	RMSE	Rank RMSE	CORR	Rank CORR	BIAS	Rank BIAS	Total Ranking
Case 1: 29 May 11 UTC - 20 UTC							
Row 1	0.811	8	0.910	1	0.708	8	17
Row 2	0.746	5	0.811	4	0.175	5	14
Row 3	0.764	6	0.799	7	0.159	4	17
Row 4	0.766	7	0.797	8	0.103	3	18
Row 5	0.684	3	0.803	6	-0.084	2	11
all	0.639	1	0.824	3	-0.041	1	5
half	0.722	4	0.807	5	0.177	6	15
Rows 1&2	0.670	2	0.878	2	0.463	7	11
Case2: 19 July 13 UTC - 20 July 08 UTC							
Row 1	1.459	8	0.669	7	0.572	8	23
Row 2	1.435	7	0.655	8	-0.064	3	18
Row 3	1.380	5	0.680	5	-0.056	2	12
Row 4	1.382	6	0.688	4	-0.106	4	14
Row 5	1.363	3	0.705	2	-0.295	6	11
all	1.300	1	0.729	1	-0.326	7	9
half	1.373	4	0.680	6	-0.043	1	11
Rows 1&2	1.350	2	0.689	3	0.278	5	10
Case 3: 26 Aug. 00 UTC - 07 UTC							
Row 1	0.790	2	0.509	1	0.269	2	5
Row 2	1.019	5	0.355	8	-0.598	4	17
Row 3	1.071	7	0.408	6	-0.707	7	20
Row 4	1.023	6	0.398	7	-0.634	5	18
Row 5	1.086	8	0.425	4	-0.764	8	20
all	1.004	4	0.471	2	-0.697	6	12
half	0.961	3	0.411	5	-0.547	3	11
Rows 1&2	0.777	1	0.461	3	-0.116	1	5
Case 4: 26 Aug. 13 UTC - 22 UTC							
Row 1	0.906	2	0.581	1	0.468	3	6
Row 2	1.026	6	0.387	8	-0.529	4	18
Row 3	1.026	7	0.434	7	-0.585	6	20
Row 4	1.006	5	0.456	6	-0.562	5	16
Row 5	1.075	8	0.493	3	-0.704	8	19
all	0.990	4	0.574	2	-0.654	7	13
half	0.946	3	0.468	5	-0.464	2	10
Rows 1&2	0.830	1	0.468	4	0.029	1	6
Case 5: 25 Oct. 16 UTC - 26 Oct. 03 UTC							
Row 1	0.921	1	0.576	1	-0.504	1	3
Row 2	1.410	3	0.346	3	-0.982	3	9
Row 3	1.634	6	0.275	6	-1.170	6	18
Row 4	1.608	5	0.249	8	-1.071	5	18
Row 5	1.722	8	0.263	7	-1.280	8	23
all	1.663	7	0.322	4	-1.277	7	18
half	1.501	4	0.304	5	-1.051	4	13
Rows 1&2	1.106	2	0.482	2	-0.715	2	6

6.1.5.2 Impact of the assimilation

Nacelle winds differ from one wind turbine to another because of wake effects as well as operational constraints. Therefore, the medians of different thinning techniques vary (Figure 6.6). In this section I analyze the impact of data assimilation in conjunction with the selected thinning techniques. In the experiments the median of the whole wind farm (ALL), the

Table 6.3: Same as Table 6.2, but for all the test cases combined.

Aggregation	RMSE	Rank RMSE	CORR	Rank CORR	BIAS	Rank BIAS	Total Ranking
Row 1	1.106	2	0.943	1	0.323	2	5
Row 2	1.226	5	0.925	6	-0.353	4	15
Row 3	1.266	7	0.922	8	-0.414	6	21
Row 4	1.253	6	0.923	7	-0.407	5	18
Row 5	1.284	8	0.929	4	-0.585	8	20
all	1.221	4	0.938	3	-0.571	7	14
half	1.204	3	0.928	5	-0.342	3	11
Rows 1&2	1.065	1	0.941	2	0.019	1	4

Table 6.4: Experiments including the thinning strategy adopted and the wind direction data source.

Experiment	Thinning strategy	Wind direction
ALL	Median of whole wind farm	yaw angle
HALF	Median of upwind half wind farm	yaw angle
ROW1	Median of upwind row	yaw angle
ROWS1&2M2	Median of first two upwind rows	43 m from M2
ROWS1&2YAW	Median of first two upwind rows	yaw angle
MADIS	only MADIS data assimilated	N.A.
CTL	no data assimilation performed	N.A.

median of the upwind half wind farm (HALF), the median of the first upwind turbine row (ROW1) and the median of the first two upwind turbine rows were assimilated. The wind direction data source for all the experiments was the yaw angle; for the first two turbine rows combined both the yaw angle (ROWS1&2YAW) and the direction measurements at M2 (ROWS1&2M2) were used to evaluate the sensitivity of different wind direction data sources. The experiments are summarized in Table 6.4.

The results show a positive impact of the assimilation of both MADIS and wind farm data beyond 24 hours (Figure 6.7). The impact of the wind farm data assimilation alone is evident in the first 5 hours for the bias (Figure 6.7 top panel, experiment ALL). The difference between the thinning techniques is most pronounced for the analysis. Beyond a 2 hour time scale, the assimilation impact of wind farm data together with MADIS data is similar to the assimilation with only MADIS data, indicating the benefit of upper air MADIS data in the assimilation. The average distribution of these data cover most of the model domain and include data upstream (Figure 6.3), which explains the long assimilation impact beyond 24 hours. Since the location of M7 and Horns Rev are only one grid point apart, with a radius of influence of 90 km this results in minor differences between the two locations. Detailed results will be explained in this section.

6.1.5.3   Assimilation impact depending on lead times up to 24h

The bias at M7 (Figure 6.7a (left)) was reduced by all the assimilation experiments compared to the CTL up to a lead time of 18 hours. Especially up to lead time 2, the assimilation of wind farm data reduced the bias also compared to assimilating only MADIS data. The

Figure 6.7: Bias (a), CRMSE (b), Rank correlation (c) and RMSE (d) of wind speed for the experiments described in Table 6.4, and for 0-24 h lead times at M7 (left) and Horns Rev (right). The reference for the metrics (i.e., the observations) for Horns Rev are the median wind speeds from ROW1.

impact of assimilating the median of the whole wind farm can even be seen up to 5 hours. The assimilation impact in bias beyond the advection timescale is suspected to be the result of the altered atmospheric state due to the data assimilation procedure. While the bias of the experiments stays stable around a value of $1.3 \mathrm{~m\,s^{-1}}$ from lead time $1 - 12$, the bias of the CTL is slightly reduced up to lead time 4 with a rapid increase around lead times $6 - 12$. The low, but increasing, bias for the first two lead times points at advection of the data assimilation impact downstream to M7. Similar conditions apply for Horns Rev (Figure 6.7a (right)).

The CRMSE, the random component of the RMSE (Figure 6.7b (left)), is reduced by the assimilation experiments compared to the CTL for the first 4 lead times and from lead time 12 onwards. However, in between these times the CTL shows slightly better scores. This and the increase of CRMSE after lead time 1 at Horns Rev compared to the CTL suggests the introduction of numerical noise through the assimilation (Figure 6.7b (right)). The reduction in CRMSE is not related to advection but due to the intrinsic skill of FDDA to improve the random error.

The RMSE (Figure 6.7d (left)) shows the added contributions of CRMSE and bias. The RMSE at M7 is reduced by the assimilation for all simulated lead times and even beyond that. The assimilation of wind farm data outperformed the MADIS experiment up to lead time 2. Experiment ALL shows the best results, and its impact is visible far beyond the advection time scale. At Horns Rev, the contribution of CRMSE to the RMSE leads to improvements up to lead time 1 and from $6 - 24$. Here, the increase in CRMSE after lead time 2 lead to a better performance of the CTL for lead times $2 - 5$.

The rank correlation (Figure 6.7c (left)) is similar for all the experiments and is nearly constant for the experiments where data assimilation was performed, whereas the correlation of the CTL drops considerably after lead time 12.

Note that Figure 6.7 shows average error metrics over all the cases. The error metrics for all the cases separately show different features (not shown). A reason for the results is thus difficult to find and drawing overall conclusions from single cases is questionable. I attempt to analyze the impact of the data assimilation further in sections 6.1.5.5 and 6.1.5.6.

6.1.5.4   Difference in thinning techniques for lead times $0 - 6\mathrm{h}$

The distinction between the thinning techniques is most pronounced during the first 5 lead times (Figure 6.7). The experiments using the yaw angle slightly outperform the ones using the wind direction measurements from M2 to the north of the farm. This is an advantage, because if the yaw angle data are known to be reliable, there is no need to additionally set up an expensive measurement mast. Moreover, only wind farm data, which are recorded for control purposes in any case, have the potential to contribute successfully to the assimilation experiments.

The verification with M7 data shows that the bias is positive, which is in agreement with the results found in Draxl et al. (2011b) for a nearby site. The bias was reduced compared to the CTL, but also compared to the run with only MADIS data (Figure 6.7 a). Experiment ALL shows the best result. The bias of ALL increases from lead time $0 - 2$, but is still the lowest during the first 6 lead times. The increase in bias up to lead time 2 for all the data
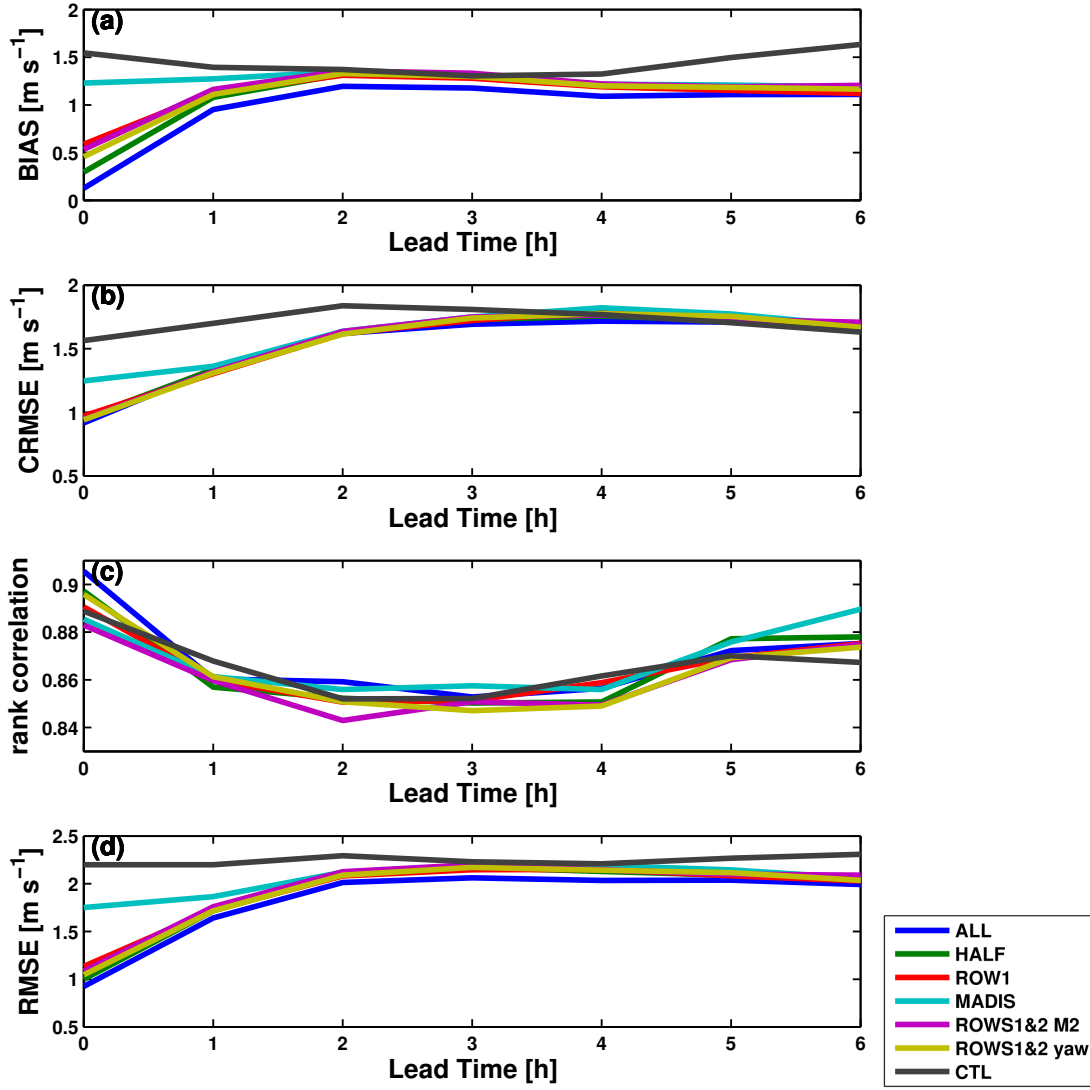
Figure 6.8: Same as Figure 6.7 (left), but for 0-6 h forecast lead times.

assimilation experiments could be due to increased noise in the simulations. For the CTL the pattern looks different: a slight reduction in bias until lead time 3, followed by a steep increase. Experiment ALL assimilated the median of all the turbines, which is lower than the one from other groups. As such it is more similar to M7. Moreover, when verifying with the wind farm itself (i.e., with the median of ROW1 for all the cases), ALL shows the best results as well during the first 5 hours (Figure 6.7a (right)), even though ROW 1 would be expected to have an advantage as the reference. This is because the assimilation decreases the wind speeds in the model estimate. ALL is showing the lowest median wind speeds (Figure 6.6), and thus most successful in doing so.

The CRMSE (Figure 6.7b), was reduced by the assimilation of wind farm data up to lead time 4 compared to experiment MADIS and the CTL, which shows the ability of the assimilation to add predictive skill to the raw forecast. One could argue that the bias of $\sim 1 \text{ m s}^{-1}$ can be the outcome of wake effects on M7. However, this effect does not explain the behavior of CRMSE. The different thinning techniques show similar behavior at the location downwind
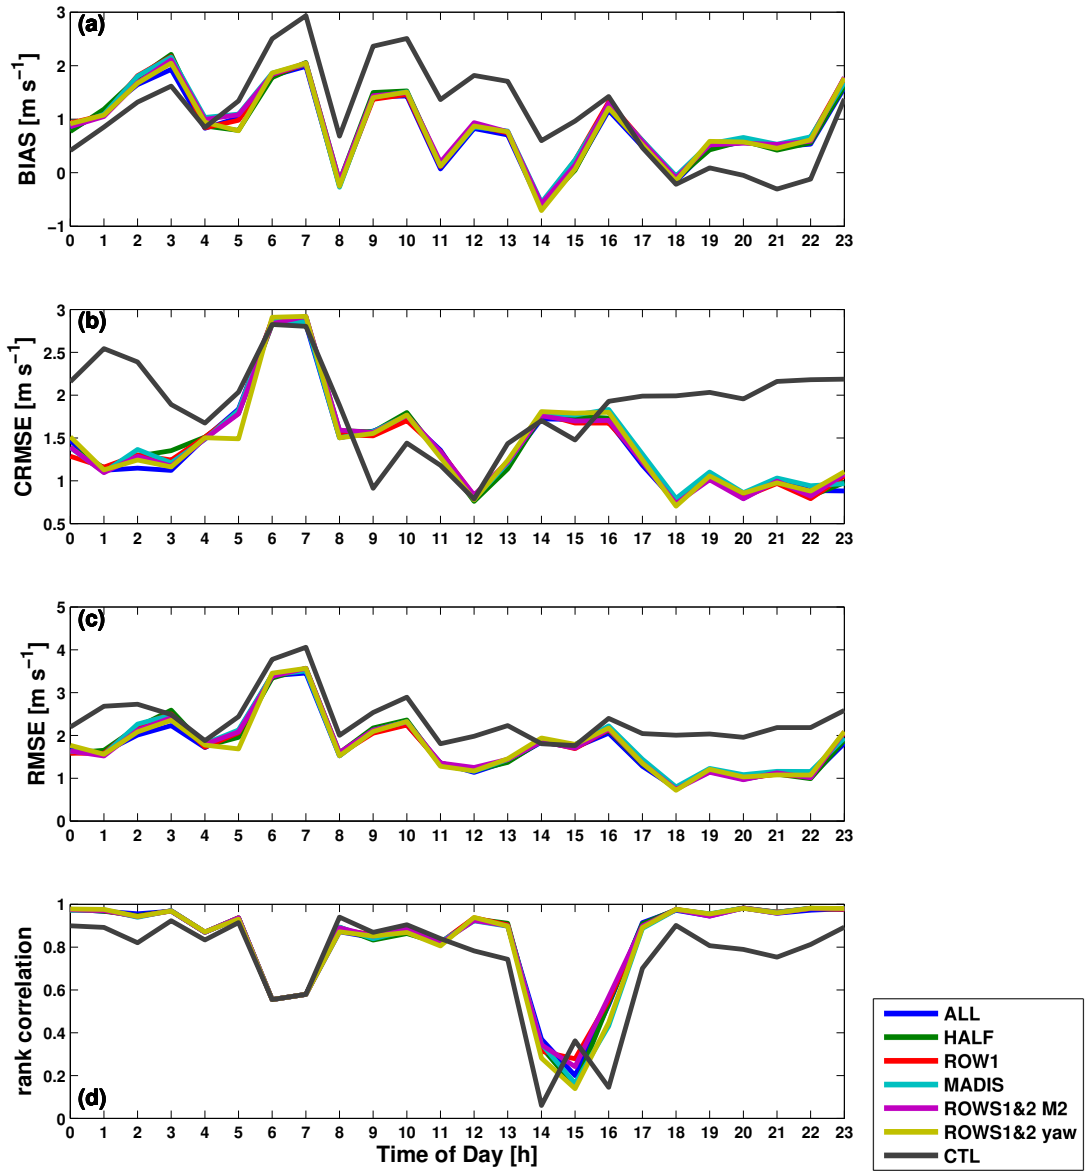
Figure 6.9: Bias (a), CRMSE (b), rank correlation (c) and RMSE (d) of wind speed for the experiments described in Table 6.4, and for each hour of the day (valid time in UTC) at M7.

and the windfarm itself; for Horns Rev, the benefit of the assimilation does not persist beyond lead time 2 for the CRMSE (Figure 6.7b (right)). The results for the bias and the CRMSE combined, shown as the RMSE in Figure 6.7d, clearly demonstrate that ALL outperforms the other experiments for the verification with M7. The rank correlation (Figures 6.7c) shows little variations ($\sim 0.1\%$) between all the experiments for lead times 0 – 6.

6.1.5.5   Assimilation impact depending on time of day

The results in the figures discussed above show the assimilation impact on the forecast lead time, but do not give information about the dependence of the results on the valid time, i.e.,

time of the day. Bias, CRMSE, rank correlation and RMSE of M7 were thus averaged over their valid time in Figure 6.9. Due to the relatively small sample size during each time of the day, the curves are uneven. However, the CTL and runs where data assimilation was performed are distinct. The bias and CRMSE show a mirrowed behaviour: whereas the systematic error was improved during the day between 5 and 16 UTC and declined during the night between 16 and 4 UTC by $0.5 \, \mathrm{m\,s^{-1}}$, the CRMSE was improved during the night between 16 and 6 UTC by $\sim 1 \, \mathrm{m\,s^{-1}}$. During the day the simulations with and without data assimilation perform similarly for CRMSE. The times of transition between improvement and decline of the error metrics in the afternoon coincides approximately with the transition from day time to night time conditions in the atmosphere (convective and stable boundary layers, respectively). The rank correlation was improved during the first lead times up to 5 UTC and after 11 UTC. Both the assimilation experiments and CTL show a deep drop in correlation around $14 - 16$ UTC. This indicates that the drop in rank correlation in Figure 6.7 was related to the valid time rather than on lead time. In fact, the drop in rank correlation in Figure 6.7 was averaged out in the assimilation experiments. The behaviour in bias and CRMSE combined is reflected in Figure 6.9d, which shows an improvement of the runs where data assimilation was performed throughout the whole day, except for $3 - 4$ and $14 - 15$ UTC. The behavior of all the assimilation experiments is similar.

6.1.5.6 Breakdown of forecast error: dependence on lead time

Figure 6.10 shows a break down of the forecast error for all the simulated lead times for the CTL and experiments ALL, HALF and ROW1 to have an overview of the assimilation impact. The forecast error is shown as a function of data assimilation experiment (*Run*, x-axis) and lead time (*hours*, y-axis), thus every colored square represents the forecast error at a particular forecast hour. The black vertical lines separate the test cases from each other. All the experiments show similar results with minor differences for specific hours (not shown for ROWS1&2 and MADIS). The experiments succeded in decreasing the forecast error in many cases, which is indicated by more blueish/greenish colors in Figures 6.10 (upper right and lower panels) compared to the CTL (Figure 6.10 upper left). This is in agreement with the results in Figure 6.7a(left), where the positive bias was decreased. The diagonal patterns suggests that the forecast error does not primarily depend on a specific lead time, but on a specific valid time. However, since the diagonal patterns are not consistent throughout the cases, the error depends rather on a specific weather situation at a certain time of the day. This is confirmed by Figure 6.11, which indicates the dependence of the error on time of the day for all the experiments. Moreover, the CTL shows a stronger and more consistent line pattern than the assimilation experiments, indicating stronger dependence on weather situation.

The diagonal patterns in Figure 6.11 (middle panel) for the CTL for e.g., runs $9 - 18$ and valid times $14 - 23$ are not clearly visible in the assimilation runs (Figure 6.11 top panel). Diagonal signals indicate a dependence on lead time; in this specific example the analysis (lead time 0) in the CTL has a different error behavior than the following lead times. Since that clear signal of diagonal line is only slightly visible in the assimilation runs, this indicates that in the assimilation experiments especially the error of the analysis was improved. At the same time, since the diagonal pattern in this example is not continuous, there is an impact of weather situation discernible as well.

Figure 6.10: Break down of the wind speed error (model minus observation) as a function of data assimilation experiment (Run; x-axis) and lead time (y-axis) for the CTL (left) and experiments ALL (upper right), HALF (lower left), ROW1 (lower right) at M7. Every colored square represents the forecast error at a particular forecast hour. The black vertical lines separate the test cases from each other. Data void areas are left white.

The dependence on the weather situation can further be seen in Figure 6.11 (lower panel), which shows the absolute differences of errors (the errors of the assimilation experiments minus the error of the CTL). While the differences between the errors of the CTL and the assimilation experiments are mostly very low (blue color), 3 blocks of higher differences stand out: The red triangle shaped blocks during runs 15 – 24 and the squared greenish blue block during runs 25 – 29. Although the performance of the experiments compared to observations was not evaluated in this figure, it shows very well the impact of the assimilation in terms of difference to the CTL and indicates weather dependence once again. During case 2 and 3 the differences were most pronounced. During case 2 the CTL exhibited a negative bias during the night, which could be successfully decreased by the assimilation. The absolute error difference during case 3 is due to a positive model bias around noon, which was reduced by the assimilation to a negative bias.

Figure 6.11: Break down of the wind speed error (model minus observation) as a function of forecast run (x-axis) and valid time (y-axis) for experiments ALL (upper left), HALF (upper middle), ROW1 (upper right) and the CTL (middle) at M7. The differences of forecast errors of the experiments to the CTL are shown for experiments ALL (lower left), HALF (lower middle), ROW1 (lower right). The black vertical lines separate the test cases from each other. Data void areas are kept white.

6.1.6 Discussion

This section explores the aspects and benefits of the assimilation of wind farm data, i.e., nacelle winds and yaw angles, with the WRF Four-Dimensional Data Assimilation system. The challenges associated with these data including wake effects on downwind turbines and disturbances of the measurements caused by the rotor blades are discussed. The focus is on the evaluation of different thinning strategies of wind farm data. Results reveal that using the median wind speed of the whole wind farm performs best compared to the other thinning strategies, which included the median of the upwind turbine row, the median of the first two upwind turbine rows, and the median of the upwind half of the wind farm.

Adding wind farm observations to already existing data used operationally to improve the initial conditions of a NWP model has the potential to produce more accurate low level wind predictions for wind energy applications both downstream of the wind farm and on the wind

farm where data were assimilated. These results are relevant for the wind energy industry: Firstly, improving wind forecasts downstream of a wind farm is especially beneficial in areas with a high number of wind farms (e.g., the North Sea), where the assimilation of wind farm data from one farm has the potential to improve wind predictions for a nearby wind farm. Furthermore, since offshore wind farms are generally near the coast, assimilating their data has the potential to improve forecasts inland. Secondly, improving the accuracy of wind forecasts translates into a huge cost benefit for wind farm operators; since the power in the wind depends on the wind speed cubed, improved wind forecasts even for a short temporal scale and magnitude result in savings for Transmission System Operators and wind farm operators. Moreover, single poorly forecast events contribute heavily to the overall cost. In this study test cases were thus evaluated both separately, which gives an indication about the assimilation impact depending on weather situation, as well as through error metrics, which describe the overall assimilation impact and include outliers in their performance (i.e., RMSE, bias, CRMSE and rank correlation). The latter help assessing the overall benefit for wind energy applications, which was the aim in this study.

Initial computations of error metrics carried out with nacelle winds and wind measurements from M7 revealed that the median of the first two upwind turbine rows would have the lowest combined error metric. However, using the median wind speed of the whole wind farm performed best in the data assimilation experiments compared to the other thinning strategies. Since the model bias is positive, the most successful thinning strategy is the one that nudges the lowest wind speeds. The median of the whole wind farm was mostly lower than the other strategies and thus the most successful. This implies, that the best thinning strategy could be different when other NWP models are used or if the model behaviour at the location of interest is different from the one used in this study. Taking the median of the whole wind farm is very practical from an operational point of view, because it does not require a real-time pre-processing algorithm to detect which row or turbine is upwind. However, down regulations, maintainance periods, and outages will decrease data quality and will have to be considered when assimilating wind farm data operationally. Wind direction measurements were not sensitive to the results; yaw angles constitute thus a potential data set that can be used successfully in data assimilation. This has the advantage that wind farm data can be used without additional (and usually expensive) tower measurements.

Moreover, transmission system operators usually convert the average wind forecast at hub height to one power forecast for the whole wind farm. When this wind forecast was initialized by assimilating the median of all turbines, which encompasses both wake affected and non-affected wind speeds and is thus lower, more realistic power estimates can be expected. True power distributions within a wind farm are also driven by wake affected and non-affected wind speeds. This was also concluded in Cutler et al. (2011).

The assimilation impact lasted beyond 24 hours, which is far beyond the advection timescale. The impact beyond ~ 2 hours was mainly due to the additional assimilation of upper air MADIS data, whereas the assimilation of wind farm data showed skill in improving CRMSE and RMSE for the first ~ 2 free forecast hours, and in bias for up to 5 hours. This is within the timescale where the forecast improvement for wind energy applications is most valuable (Chapter 2). The improvement of all these metrics together implies that the positive impact of wind farm data assimilation was not solely due to their advection downwind, but due the ability of FDDA to add predictive skill to the raw forecast; moreover, the fact that M7 and Horns Rev are both within the radius of influence of the wind farm data assimilation contributed further to the improvement of these metrics.

Wind farm data were assimilated together with upper air MADIS data, because they constitute a standard set of measurements in current operational forecasting systems, and because more advanced data assimilation techniques require upper air observations for the data assimilation effect to stay in the system (section 4.2). Independent simulations assimilating wind farm data only into FDDA for another day and different weather situation showed similar results to the ones presented here: The impact of the assimilation lasted for 2 hours.

An improved CRMSE suggests a reduction of phase errors. The rank correlation was mostly improved for lead times beyond 13 hours, which was not due to the assimilation of wind farm data. The difference in thinning strategies is most significant for the first 5 lead times. The assimilation impact depends on the time of the day, the forecast lead time and the weather situation. Even though all the test cases occurred during similar weather situations with unstable atmospheric conditions influenced by low pressure systems (section 6.1.3, Figure 6.4), different assimilation effects were encountered. It is likely that the starting time of the simulations influenced the results as well. Due to the small sample size this effect was not quantifiable.

Most of the assimilation impact in the model domain is due to the additionally assimilated upper air observations. However, the additional impact of wind farm data is considerable. Figure 6.12 shows wind speed increments for the first free forecast hour of experiment ALL for four selected examples during different times of the day. The increments compared to the CTL are shown in the first column. The contribution of wind farm data in the assimilation process can be seen in the second column, which indicates that the biggest impact is not necessarily close to the wind farm but can be a few hundred kilometers away and depends on the case. In areas sensitive to the assimilation, employing different thinning techniques - shown here differences ALL minus ROW1 - triggers wind field differences throughout the domain (third column). This analysis clearly shows that the biggest impact of the assimilation is not necessarily by or downstream of the wind farm and explains assimilation impacts beyond the first few forecast lead times due to an impact accross the whole model domain, that is not necessarily where data have been assimilated.

The restrictions for the case studies, although necessary for the conducted research, resulted in a limited number of cases and the results are thus likely to not be statistically significant. However, results are relevant for the wind energy community and show a clear trend. Thus, for future studies a setup is recommended that allows case studies that yield more statistically robust data. This will make it necessary to re-define the "first upwind row", which is straightforward in regularly shaped wind farms, but gets more complicated when the wind is not coming directly perpendicular towards a regularly shaped wind farm edge. Moreover, wakes depend besides wind speeds, atmospheric stability and turbulence intensity on wind farm layouts and farm size. The best thinning strategies for irregularly shaped and bigger wind farms might therefore be different and will have to be re-assessed.

The methods and issues concerning the assimilation of wind farm data on the example of FDDA are valid for any other data assimilation system alike. The benefit of assimilating wind farm observations in these other systems will be addressed in the future. Likewise, the assimilation of observations in the boundary layer is a topic that still needs further research.

Since the field of wind energy is a commercial enterprise, wind farm operators are reluctant to provide their data. To assimilate wind farm data in an operational context, lobbying work will be necessary to make the data widely accessible. Hopefully this study is encouraging to
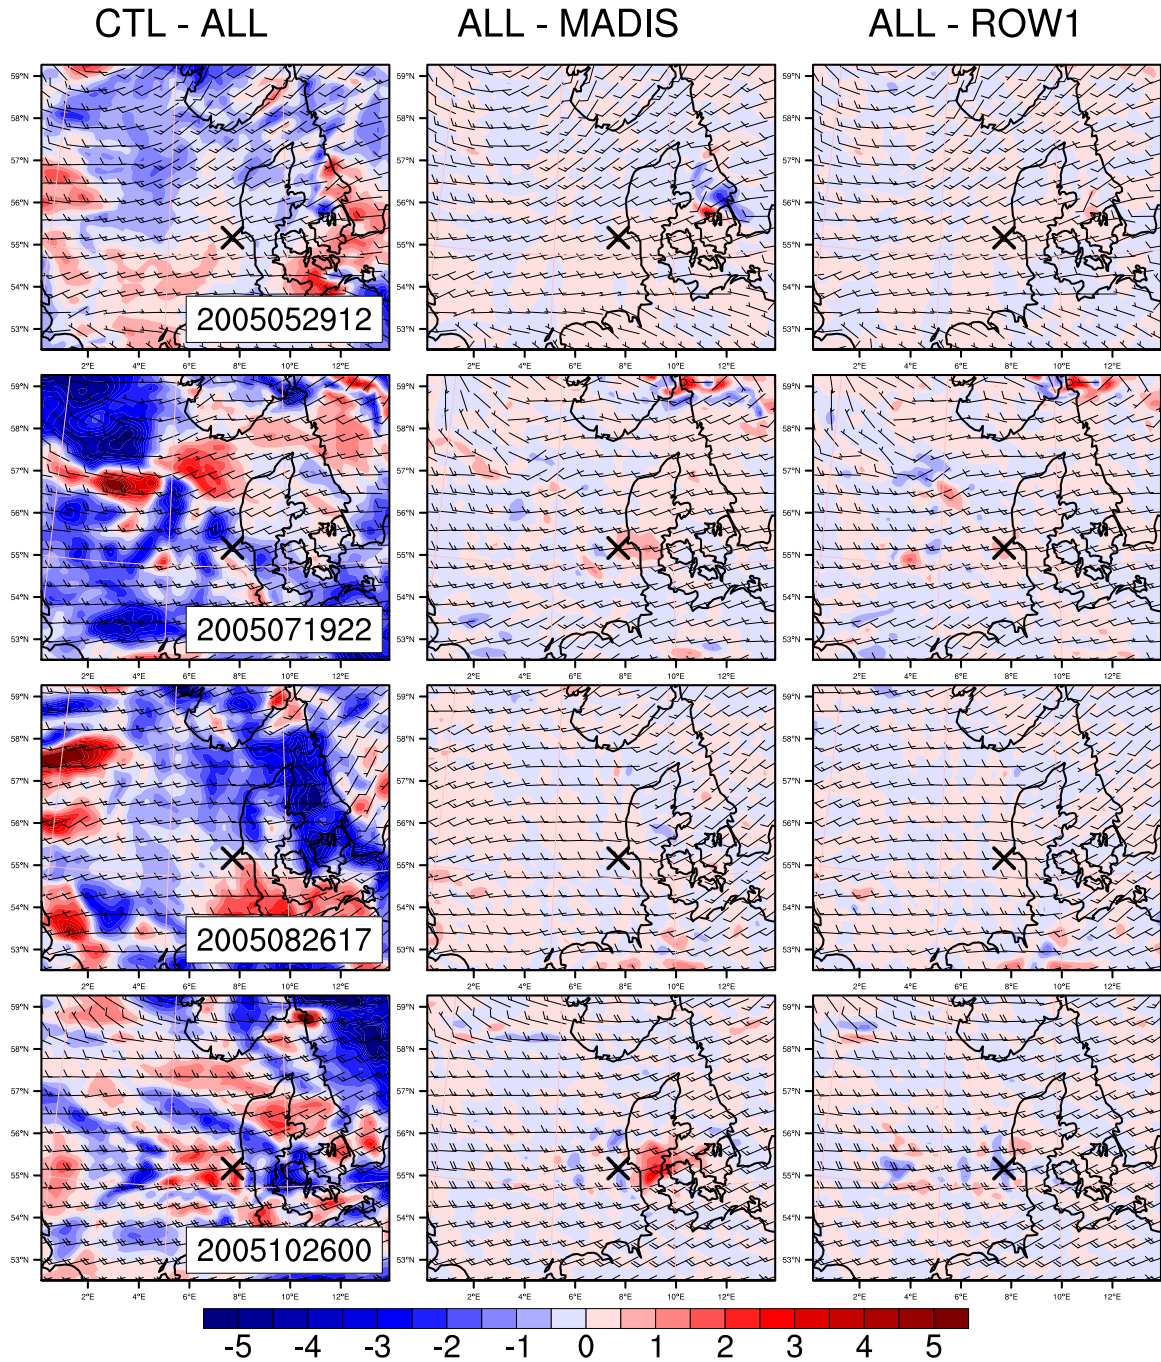
Figure 6.12: Wind speed increments in $\mathrm{m\,s^{-1}}$ at the third model level (around 117 m agl) for experiments CTL minus ALL (first column), ALL minus MADIS (second column) and ALL - ROW1 (last column) for 29 May 12 UTC (first row), 19 July 22 UTC (second row), 26 August 17 UTC (third row) and 26 October 00 UTC (last row), which was the time of the first free forecast hour. The cross shows the location of Horns Rev and the wind barbs represent the wind from the CTL in column 1 and from experiment ALL in columns 2 and 3.

advance initial steps, and this rich new data set will be widely used in data assimilation in the future.

6.2 Qualitative assessment of the radius of influence

The FDDA system has not been tuned for the above study. In order to analyze models for their relative performance, tuning is not necessary and results will not be very different. However, for an analysis of the absolute performance of a data assimilation system, tuning will lead to improved forecast statistics. This section will thus qualitatively assess the impact of the radius of influence (section 4.1.1 and 4.1.2).

The weather situation plays a role on whether the assimilation impact is advected away or wiped out or stays in the system. The day of research for this assessment is the 24 June 2009. It is characterised by a high pressure system residing over Scandinavia. The center of the high was over the North Sea on the 24th. Denmark was constantly in the south of the center of the high with relatively light winds. Forcings were local which is advantageous for the data assimilation impact to stay in the system.

The experiments start assimilating data on the 24 June 2009 at 00 UTC, assimilating each hour for a period for 6 hours. After that the free forecast starts. The domain consisted of 3 nested domains similar to Figure 6.1, but additionally with a third domain with a grid spacing of 3.3 km. All the domains used 57 model levels, with 8 levels within the lowest 500 m. The lowest levels which are important for wind energy applications were located at approximately 12, 46, 96, 156 m. .

Any assimilation is supposed to impact the original wind field horizontally and vertically. To which extent depends on the assimilation method (section 4.1.2) and other tuning parameters, one of which is the radius of influence of an observation. In flat terrain like the area around Horns Rev, increasing the radius of influence should have an impact that reaches further, which was confirmed in Figure 6.13. Wind speed increments are shown for the asimilation of nacelle winds for a radius of influence of a) 120 km, b) 80 km, c) 40 km and d) 20 km for 06 UTC, after 6 hours of nudging. The larger the radius of influence, the bigger the impact in the analysis. Note that the larger/smaller the impact near the wind farm, the smaller/larger the impact on the mainland of Denmark. This might be due to geostrophical balancing, which was more pronounced when the radius of influence was smaller. Note that 40 km was the radius of influence used in the assimilation experiments in section 6.1. Wind speed increments also differ with height and the assimilation of only one observation with different radii of influence shows differences up to 2400 m agl (Figure 6.14). Unfortunately no observations were available to verify the model runs up to this height. However, this qualitative assessment shows the sensitivity of observation weighting to model results.

6.3 Conclusions

This chapter presents a first assessment of the assimilation of wind farm data with relevant results for the wind energy industry: they show that nacelle winds and yaw angles from wind farms are a promising data set to improve wind predictions both downstream and on the wind farm where data were assimilated. This is especially beneficial in areas with a high number of wind farms (e.g., the North Sea), where the assimilation of wind farm data from one farm

Figure 6.13: Wind speed increments in m s^{-1} at the second model level ($\sim$46 m agl) for the assimilation of nacelle winds at the location of the wind farm Horns Rev I (denoted by the cross) for a radius of influence of (a) 120 km, (b) 80 km, (c) 40 km and (d) 20 km after 6 hours of nudging. The wind barbs are taken from the original forecast (first guess).

has the potential to improve wind predictions for a nearby wind farm. Moreover, having additional data is beneficial in areas with low observational coverage. Furthermore, since offshore wind farms are generally near the coast, assimilating their data has the potential to improve forecasts inland (as could be confirmed by Figure 6.12). With more and more wind farms being built and wind farm observations becoming available, adding these to already existing data used operationally to improve the initial conditions of a NWP model will be promising for both wind energy predictions and weather predictions in general.

Results revealed that using the median wind speed of the whole wind farm performed best compared to the other thinning strategies tested in this study. It was further shown, that yaw angles constitute a potential data set that can be used successfully in data assimilation. This has the advantage that wind farm data can be used without additional (and usually expensive) tower measurements.

The methods and issues concerning the assimilation of wind farm data on the example of FDDA are valid for any other data assimilation system alike. However, how to assimilate observations in the boundary layer is a topic that still needs further research in more advanced data assimilation systems (section 4.2).
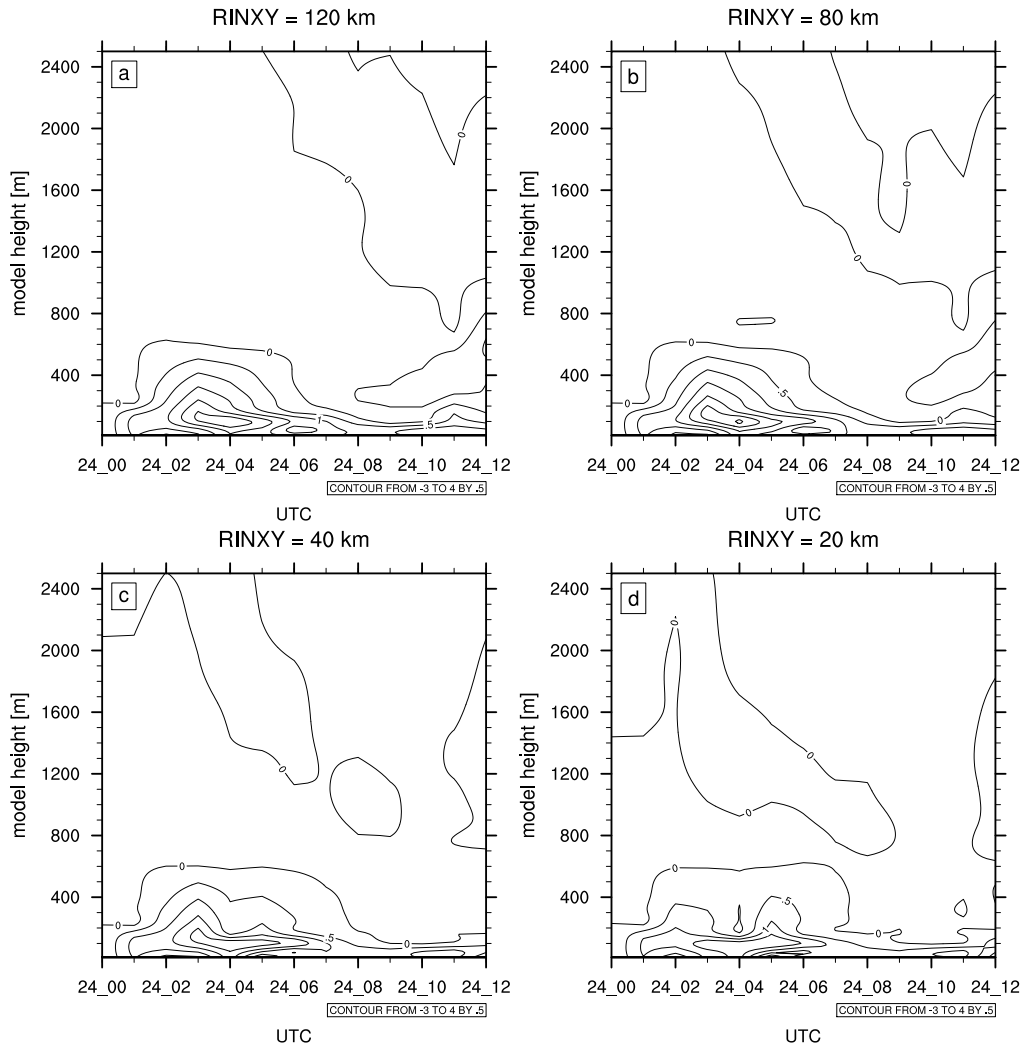
Figure 6.14: Wind speed increments as a function of height (x-axis) and time (y-axis) in $\mathrm{m\,s^{-1}}$ for the grid point closest to Horns Rev I for a radius of influence of (a) 120 km, (b) 80 km, (c) 40 km and (d) 20 km. Nacelle winds were nudged for 6 hours (until 24 June 06 UTC), then the free forecast starts.

CHAPTER 7

Summary and Conclusions

The introduction of significant amounts of wind energy into power systems necessitates accurate wind forecasts at wind farm locations. Since the power in the wind is proportional to the wind speed cubed even small wind forecast errors result in large power prediction errors. In this context wind energy forecasting relates to forecasting wind speeds because it is the main quantity needed and commonly used to predict power output.

Accurate wind forecasts are worth billions of dollars annually; forecast improvements will result in reduced costs to consumers due to better integration of wind power into the power grid, and more efficient trading of wind power on energy markets. Additionally, improved wind forecasts will enhance the value and acceptance of wind power, increase the share of renewable energies in the power mix and as such help reduce dependence on fossil fuels. That will be to the benefit to all of us by mitigating the interference of fossil fuels with climate and environment.

The predictability of hub height winds has been examined with the WRF model. Although wind speeds at hub height are usually used to predict power output, it was shown that taking hub height winds only is not sufficient. Calculating the power output from hub height winds alone may result in erroneous estimates due to the vertical wind shear in the PBL. Therefore, a study was undertaken to analyze the results of numerical experiments with the WRF model using seven different PBL parameterizations regarding their performance in forecasting wind speeds for wind energy applications. The accuracy of modeled wind conditions in the PBL depends, among several modeling aspects, on the PBL scheme adopted. The ability of the PBL schemes to forecast the mean wind speed, its time variability and wind shear depends on atmospheric static stability. The MYJ parameterization performs best during stable and very stable atmospheric conditions and the ACM2 during neutral and near stable cases. The YSU scheme outperforms the others during unstable conditions.

The choice of PBL scheme for a forecasting system for a particular region will thus depend on the typical distribution of atmospheric stability conditions at the site. The results further signal that moving beyond the paradigm of 10 m and hub height wind verification alone is essential for wind energy applications. More truthful results can be expected for site

assessments or wind energy forecasts, when a NWP model is validated at different heights across the rotor area. This is important for wind energy applications, especially offshore where many of the new wind farms will be built, since shallow stable boundary layers occur frequently and can result in excessive shear. In fact, a novel aspect to this study is the presentation of a verification methodology that takes into account wind at different heights where turbines operate.

Improving the model's initial conditions through data assimilation to yield more accurate wind predictions was also shown to be a crucial component of wind energy forecasting. With an increasing number of wind farms being built, a new and unique data set is becoming available as a beneficial addition to already existing data sets used in data assimilation: wind speeds measured on the nacelle of a wind turbine and the yaw angle, the turbines angle of rotation into the wind. A positive feedback loop can result from building additional wind farms: each additional wind farm's observations can be assimilated to improve other wind forecasts.

Since the observations of nacelle winds on the individual wind turbines are not independent from each other and dealing with each turbine's measurements separately makes the data handling challenging, thinning the data is desirable. A study of thinning approaches was undertaken to explore different thinning strategies and their impact on four-dimensional data assimilation concerning wind predictions at hub height for the wind farm Horns Rev I. Results revealed that using the median wind speed of the whole wind farm performed best compared to other thinning strategies. This is very practical from an operational point of view, since this strategy does not require a real-time pre-processing algorithm to detect which row or turbine is upwind. The new observations are of significant value in improving wind forecasts downstream and on the wind farm itself for the first five forecast hours. Since the value of an improved forecast is particularly high for a forecast horizon of up to a few hours, these results are very signicant. Moreover, it was shown that the improvement of the forecasts was not solely due to the advection of the assimilation impact downstream, but due to the ability of FDDA to add predictive skill to the raw forecast. The assimilation impact depends on the time of the day, the forecast lead time and the weather situation. The results from this study are especially promising for areas with a high number of wind farms (e.g., the North Sea), where the assimilation of wind farm data from one farm has the potential to improve wind predictions for a nearby farm. Furthermore, since offshore wind farms are generally near the coast, assimilating their data has the potential to improve inland forecasts as well. This thesis is a first step towards the endeavour to allow wind farm data to be considered potential candidates in contributing to a rich data set in the PBL for data assimilation.

CHAPTER 8

Implications and Future Work

While this thesis has addressed limitations of wind energy forecasting with mesoscale models and techniques for improvement, it would be very desirable to end users that the results be incorporated into an operational framework. This would require that the NWP model be tuned to transmission system operations needs, and to the conditions at wind farm sites and the forecast variables of interest in order to improve the absolute forecast statistics.

In section 3.3 the focus was on simulating wind and its vertical structure using seven WRF PBL parameterizations. Analyses concerning model behavior in predicting eddy diffusivity, TKE, Prandtl number, mixing length, virtual potential temperature profiles and boundary layer heights can lead to more explanations on the different behavior of the schemes. This would be a suggestion for further studies and was partly conducted for the YSU, ACM2, MYJ, and QNSE PBL parameterizations in Shin and Hong (2011). Furthermore, since a PBL parameterization is not applied in isolation from the other settings of the model, interesting future work would be to explore other aspects of the model within the verification framework that is proposed in this paper. Evaluating the forecasts at different sites could lead to interesting results as well. A big shortcoming is the lack of tall measurement masts and the low availability of flux measurements at most sites.

Although the results in this thesis are derived from model simulations covering a brief period, they mirror the behavior of the model very well. Long term operational setups would lead to more statistically robust results. However, further shorter and less consistent results similiar to those from Chapter 3 and Chapter 6 have shown that the conclusions are robust. For the assimilation of wind farm data, further studies about non-symmetrically shaped wind farms, different weather situations over a few months or on-shore wind farms are likely to indicate further relevant conclusions for wind energy predictions. It should be noted that the assimilation of observations from wind farms onshore or in complex terrain may lead to different results than the ones presented here for the offshore wind farm Horns Rev I.

When dealing with the assimilation of observations close to the surface, data assimilation quality control algorithms often reject data due to the mismatch of the model's topography and the actual physical topography in the real world. This is mainly caused by the limited

resolution of a mesoscale model. Therefore, when dealing with wind farm data in complex terrain, topographical effects will also have to be addressed in addition to the disturbances caused by turbine rotors. This is also the case with tower data used for verification when wind observations are compared to winds based on a mesoscale model. A post-processing procedure based on WAsP which is able to use output from several mesoscale models (KAMM, WRF, AROME) and to extend mesoscale model output to much finer resolutions could be used for these purposes (Badger et al., 2010).

Due to the lack of suitable tower or sodar/lidar measurements during the data assimilation experiments, the impact of the assimilation of wind farm data could not be studied within the range of a turbine height, i.e., 30 – 120 m AGL. For future studies it will be crucial to use the verification techniques in Chapter 3 that take different heights across the turbine rotor into account. This might reveal interesting insight into the impact of the assimilation of wind farm observations. Observations that cover the whole rotor area will have to be available for that purpose.

The verification of the data assimilation impact on other wind farms nearby, on their combined power output, and on the impact on bidding areas or geographical areas of transmission system operators will be relevant for future studies. This is due to the fact that the financial benefit is closely related to the combined power output of such an area. For such a study, a mechanism to convert wind forecasts to the power in a wind farm is required, which takes the wind distribution with height and wake effects into account. Due to the number of factors in energy consumption and market behavior, an estimate of the financial benefit of an improved forecast is very challenging if not impossible to calculate (Marquis et al., 2011). This information would yet be very valuable for day-to-day operations, long-term investment planning, policy-making and the overall acceptance of wind power.

The data assimilation experiments in this thesis were based on FDDA. For future studies of nacelle wind speeds and yaw angles in data assimilation systems other than FDDA it would be favorable to develop an observation operator which will account for the data quality issues associated with wind farm data and derive representative information. Developing an observation operator using power data directly is more challenging as discussed in section 5.5.

Exploring the benefit of wind farm data in advanced data assimilation techniques may increase the accuracy of wind forecasts for wind energy purposes and forecasts as a whole. Ensemble Kalman Filters are appealing in this case, because they provide flow-dependent error statistics and therefore are expected to produce an analysis with smaller errors than schemes employing fixed-error statistics (sections 4.1.3 and 4.2). As further elaborated in section 4.2, the important questions of how to most effectively assimilate PBL wind observations are still open. Since the accuracy of modeled wind conditions and wind profiles in the PBL depends, among several modeling aspects, on the PBL scheme adopted and differs under varying atmospheric stability conditions (Draxl et al. (2011b) and Chapter 3), accounting for atmospheric stability effects in data assimilation could be one way of addressing the issue. The EnKF has a strong dependence on the accuracy of the model, and in the PBL a significant portion of model error is due to uncertainties in PBL parameterizations (section 4.2). Future research regarding EnKF could thus be dealing with covariance inflation procedures, bias correction algorithms, stochastic approaches, or the use of multi-physics ensembles. One other reason why EnKF have issues in the PBL is due to sampling error. Localization is a means to ameliorate sampling error when small ensembles are used. Future

research could therefore be on the area of influence of an observation which should change with atmospheric stability. Questions to be addressed in an EnKF system could thus be: How should atmospheric stability be accounted for in ensemble data assimilation systems? How does atmospheric stability affect forecast error covariance? How should flow-dependent covariance localization in different atmospheric stability conditions be addressed? Such a study would contribute to the understanding of data assimilation in the PBL and results could be applied to the assimilation of surface observations in general. Moreover, such results have the potential to lead to improvements in the PBL/surface layer parameterizations in WRF and other numerical weather predictions systems.

Bibliography

Akhamatov, V., C. Rasmussen, P. Eriksen, and J. Pedersen, 2007: Technical aspects of status and expected future trends for wind power in Denmark. *Wind Energy*, **10**, 31–49.

Alapaty, K., N. L. Seaman, D. S. Niyogi, and A. F. Hanna, 2001: Assimilating surface data to improve the accuracy of atmospheric boundary layer simulations. *J. Appl. Meteor.*, **40**, 2068–2082.

Anderson, J., T. Hoar, K. Raeder, H. Liu, N. Collins, R. Torn, and A. Arellano, 2009: The data assimilation research testbed: A community data assimilation facility. *Bulletin of the American Meteorological Society*, **90**, 1283–1296, doi:10.1175/2009BAMS2618.1.

Anderson, J. L., 2007a: An adaptive covariance inflation error correction algorithm for ensemble filters. *Tellus A*, **59**, 210–224, doi:10.1111/j.1600-0870.2006.00216.x.

— 2007b: Exploring the need for localization in ensemble data assimilation using a hierarchical ensemble filter. *Physica D: Nonlinear Phenomena*, **230**, 99–111.

Anthes, R. A., 1974: Data assimilation and initialization of hurricane prediction models. *J. Atmos. Sci.*, **31**, 702–719.

Antoniou, I. and T. F. Pedersen, 1997: Nacelle anemometry on a 1 MW wind turbine. Technical Report Risø-R-941(EN), Risø DTU.

Badger, J., X. G. Larsen, N. Mortensen, A. N. Hahmann, J. Hansen, and H. E. Jørgensen, 2010: A universal mesoscale to microscale modelling interface tool. *European Wind Energy Conference*, Warsaw, Poland.

Bannister, R. N., 2008: A review of forecast error covariance statistics in atmospheric variational data assimilation. I: Characteristics and measurements of forecast error covariances. *Quarterly Journal of the Royal Meteorological Society*, **134**, 1951–1970, doi:10.1002/qj.339.

Barker, D. M., W. Huang, Y.-R. Guo, A. J. Bourgeois, and Q. N. Xiao, 2004: A three-dimensional variational data assimilation system for MM5: Implementation and initial results. *Mon. Wea. Rev.*, **132**, 897–914.

Barthelmie, R. J. and L. E. Jensen, 2010: Evaluation of wind farm efficiency and wind turbine wakes at the Nysted offshore wind farm. *Wind Energy*, **13**, 573–586, doi:10.1002/we.408.

Barthelmie, R. J., S. Pryor, S. Frandsen, K. Hansen, J. G. Schepers, K. Rados, W. Schletz, A. Neubert, L. Jensen, and S. Neckelmann, 2010: Quantifying the impact of wind turbine wakes on power output at offshore wind farms. *Journal of Atmospheric and Oceanic Technology*, doi:10.1175/2010JTECHA1398.1.

Barwell, B. R. and A. C. Lorenc, 1985: A study of the impact of aircraft wind observations on a large-scale analysis and numerical weather prediction system. *Quarterly Journal of the Royal Meteorological Society*, **111**, 103–129, doi:10.1002/qj.49711146704.

Benjamin, S. G., B. D. Jamison, W. R. Moninger, S. R. Sahm, B. E. Schwartz, and T. W. Schlatter, 2009: Relative short-range forecast impact from aircraft, profiler, radiosonde, VAD, GPS-PW, METAR, and mesonet observations via the RUC hourly assimilation cycle. *Mon. Wea. Rev.*, **138**, 1319–1343, doi:10.1175/2009MWR3097.1.

Benjamin, S. O. and N. L. Seaman, 1985: A simple scheme for objective analysis in curved flow. *Mon. Wea. Rev.*, **113**, 1184–1198.

Beyer, H., T. Pahlke, W. Schmidt, H.-P. Waldl, and U. de Witt, 1994: Wake effects in a linear wind farm. *Wind Eng. Ind. Aerodyn*, **51**, 303–318.

Blumen, W., R. Banta, S. Burns, D. Fritts, R. Newsome, G. Poulos, and J. Sun, 2001: Turbulence statistics of a Kelvin–Helmholtz billow event observed in the night-time boundary layer during the Cooperative Atmosphere–Surface Exchange Study field program. *Dynamics of Atmospheres and Oceans*, **34**, 189–204.

Bosveld, F., d. Bruijn, E.I.F., and A. Holtslag, 2008: Intercomparison of single-column models for GABLS3: Preliminary results. *18th AMS workshop on Boundary Layer and Turbulence*, Stockholm, Sweden.

Botterud, A., J. Wang, V. Miranda, and R. J. Bessa, 2010: Wind power forecasting in U.S. electricity markets. *The Electricity Journal*, **23**, 71–82.

Botterud, A., J. Wang, C. Monteiro, and V. Miranda, 2009: Wind power forecasting and electricity market operations. *32nd IAEE International Conference*, San Francisco, U.S.A.

Cheng, W. Y., Y. Liu, Y. Liu, B. Mahoney, M. Politovich, T. T. Warner, K. Parks, and J. Himelic, 2011: Improving the 0-6 hour wind forecast through wind farm data assimilation in the NCAR/ATEC WRF-ARW RTFDDA. *Poster at the 12th Annual WRF Users' Workshop*, Boulder, U.S.A.

Childs, P. P., A. L. Qureshi, S. Raman, K. Alapaty, R. Ellis, R. Boyles, and D. Niyogi, 2006: Simulation of convective initiation during IHOP_2002 using the flux-adjusting surface data assimilation system (FASDAS). *Mon. Wea. Rev.*, **134**, 134–148, doi:10.1175/MWR3064.1.

Christiansen, M. B. and C. B. Hasager, 2005a: Wake effects of large offshore wind farms identified from satellite SAR. *Remote Sens. of Environ.*, **98**, 251–268.

— 2005b: Wake effects of large offshore wind farms identified from satellite SAR. *Remote Sensing of Environment*, **98**, 251 – 268, doi:10.1016/j.rse.2005.07.009.

Costa, A., A. Crespo, J. Navarro, G. Lizcano, H. Madsen, and E. Feitosa, 2008: A review on the young history of the wind power short-term prediction. *Renewable and Sustainable Energy Reviews*, **12**, 1725 – 1744, doi:10.1016/j.rser.2007.01.015.

Cutler, N. J., H. R. Outhred, and I. F. MacGill, 2011: Using nacelle-based wind speed observations to improve power curve modeling for wind power forecasting. *Wind Energy*, doi:10.1002/we.465.

Cuxart, J., A. Holtslag, R. Beare, E. Bazile, A. Beljaars, A. Cheng, L. Conangla, M. Ek, F. Freedman, R. Hamdi, A. Kerstein, H. Kitagawa, G. Lenderink, D. Lewellen, J. Mailhot, T. Mauritsen, V. Perov, G. Schayes, G.-J. Steeneveld, G. Svensson, P. Taylor, W. Weng, S. Wunsch, and K.-M. Xu, 2006: Single-column model intercomparison for a stably stratified atmospheric boundary layer. *Bound.-Layer Meteor.*, **118**, 273–303.

Daley, R., 1991: *Atmospheric data analysis*. Cambridge University Press, 455 pp.

Delle Monache, L., L. Glascoe, L. J., J. Mirocha, M. Simspon, and M. Singer, 2010: Ensemble-based data assimilation for wind energy predictions. *Fifth International Symposium on Computational Wind Engineering*, North Carolina, U.S.A.

Delle Monache, L., T. Nipen, Y. Liu, G. Roux, and R. Stull, 2011: Kalman filter and analog schemes to post-process numerical weather predictions. *Mon. Wea. Rev.*, **139**, 3554–3570, doi:10.1175/2011MWR3653.1.

Draxl, C., L. Delle Monache, F. Vandenberghe, Y. Liu, and A. Hahmann, 2011a: Thinning strategies for the assimilation of wind farm observations to improve wind predictions for wind energy applications. *submitted to Wind Energy*.

Draxl, C., A. N. Hahmann, A. Peña, and G. Giebel, 2011b: Evaluating winds and vertical wind shear from WRF model forecasts using seven PBL schemes. *conditionally accepted in Wind Energy*.

Draxl, C., C. L. Vincent, A. N. Hahmann, G. Giebel, and P. Pinson, 2009: Variability forecasts for wind farms using high resolution initial conditions. *8th International Workshop on Large-Scale Integration of Wind Power into Power Systema as well as Transmission Networks for Offshore Wind Power Plants*, Bremen, Germany.

Ernst, B., B. Oakleaf, M. Ahlstrom, M. Lange, C. Moehrlen, B. Lange, U. Focken, and K. Rohrig, 2007: Predicting the wind. *Power and Energy Magazine, IEEE*, **5**, 78 –89, doi:10.1109/MPE.2007.906306.

Espana, G., S. Aubrun, and P. Devinant, 2009: Is the meandering of a wind turbine wake due to atmospheric length scales? *Progress in Turbulence III*, J. Peinke, M. Oberlack, and A. Talamelli, eds., Springer Berlin Heidelberg, volume 131 of *Springer Proceedings in Physics*, 91–94, 10.1007/978-3-642-02225-8_21.

Fast, J. D., 1995: Mesoscale modeling and four-dimensional data assimilation in areas of highly complex terrain. *J. Appl. Meteor.*, **34**, 2762–2782.

Floors, R., S.-E. Gryning, A. Peña, and E. Batchvarova, 2011: Analysis of diabatic flow modification in the internal boundary layer. *Meteorologische Zeitschrift*, **20**, 649–659, doi:10.1127/0941-2948/2011/0290.

Focken, U., M. Lange, K. Mönnich, H.-P. Waldl, H. G. Beyer, and A. Luig, 2002: Short-term prediction of the aggregated power output of wind farms – a statistical analysis of the reduction of the prediction error by spatial smoothing effects. *Journal of Wind Engineering and Industrial Aerodynamics*, **90**, 231 – 246, doi:10.1016/S0167-6105(01)00222-7.

Frandsen, S., J. Sørensen, R. Mikkelsen, T. Pedersen, I. Antoniou, and K. Hansen, 2009: The generics of wind turbine nacelle anemometry. *European Wind Energy Conference*, Marseille, France.

Fujita, T., D. J. Stensrud, and D. C. Dowell, 2007: Surface data assimilation using an Ensemble Kalman Filter approach with initial condition and model physics uncertainties. *Mon. Wea. Rev.*, **135**, 1846–1868, doi:10.1175/MWR3391.1.

Gallego, C., P. Pinson, H. Madsen, A. Costa, and A. Cuerva, 2011: Influence of local wind speed and direction on wind power dynamics – Application to offshore very short-term forecasting. *Applied Energy*, **88**, 4087 – 4096.

Giebel, G., 2000: *On the benefits of distributed generation of wind energy in Europe*. Ph.D. thesis, Carl von Ossietzky Universität, Oldenburg, Germany.

Giebel, G., R. Brownsword, G. Kariniotakis, M. Denhard, and C. Draxl, 2011: The state-of-the-art in short-term prediction of wind power: A literature overview, 2nd edition. Project report for the ANEMOS.plus and SafeWind projects, Risø DTU.

Giebel, G. and S. E. Gryning, 2004: Shear and stability in high met masts, and how WAsP treats it. *EWEA Special Topic Conference "The Science of Making Torque from Wind"*, Delft, Netherlands.

Giebel, G., P. Sørensen, and H. Holttinen, 2007: Forecast error of aggregated wind power. Technical report, Tradewind Project, EWEA.

Gryning, S. E., E. Batchvarova, B. Brümmer, H. Jørgensen, and S. Larsen, 2007: On the extension of the wind profile over homogeneous terrain beyond the surface layer. *Bound.-Layer Meteor.*, **124**, 251–268.

Hacker, J. P. and D. Rostkier-Edelstein, 2007: PBL state estimation with surface observations, a column model, and an ensemble filter. *Mon. Wea. Rev.*, **135**, 2958–2972, doi:10.1175/MWR3443.1.

Hacker, J. P. and C. Snyder, 2005: Ensemble Kalman filter assimilation of fixed screen-height observations in a parameterized PBL. *Mon. Wea. Rev.*, **133**, 3260–3275, doi:10.1175/MWR3022.1.

Hahmann, A. N. and A. Peña, 2010: Validation of boundary-layer winds from WRF mesoscale forecasts over Denmark. *European Wind Energy Conference and Exhibition*, Warsaw, Poland.

Hahmann, A. N., D. Rostkier-Edelstein, T. T. Warner, F. Vandenberghe, Y. Liu, R. Babarsky, and S. P. Swerdlin, 2009: A reanalysis system for the generation of mesoscale climatographies. *J. Appl. Meteor. Climatol.*, **49**, 954–972, doi:10.1175/2009JAMC2351.1.

Hansen, K. S., R. J. Barthelmie, L. E. Jensen, and A. Sommer, 2012: The impact of turbulence intensity and atmospheric stability on power deficits due to wind turbine wakes at Horns Rev wind farm. *Wind Energy*, **15**, 183–196, doi:10.1002/we.512.

Hasager, C. B., A. Peña, T. Mikkelsen, M. Courtney, I. Antoniou, S.-E. Gryning, P. Hansen, and P. Sørensen, 2007: 12MW Horns Rev experiment. Technical report, Risø-R-1506(EN), Risø DTU.
URL http://130.226.56.153/rispubl/reports/ris-r-1506.pdf

Hayden, C. M., 1973: Experiments in the four-dimensional assimilation of Nimbus 4 SIRS data. *J. Appl. Meteor.*, **12**, 425–436.

Holttinen, H. and R. Hirvonen, 2005: *Power System Requirements for Wind Power*, John Wiley & Sons, Ltd, chapter 8. 143–167.

Hong, S.-Y., Y. Noh, and J. Dudhia, 2006: A new vertical diffusion package with an explicit treatment of entrainment processes. *Mon. Wea. Rev.*, **134**, 2318–2341.

Hu, X.-M., J. W. Nielsen-Gammon, and F. Zhang, 2010: Evaluation of three planetary boundary layer schemes in the WRF model. *J. Appl. Meteorol. Climatol.*, **49**, 1831–1844.

Janjic, Z. I., 2001: Nonsingular implementation of the Mellor-Yamada level 2.5 scheme in the NCEP Meso model. Technical report, National Centers for Environmental Prediction: Camp Springs, MD, U.S.A.

Jankov, I., W. A. Gallus, M. Segal, B. Shaw, and S. E. Koch, 2005: The impact of different WRF model physical parameterizations and their interactions on warm season MCS rainfall. *Wea. Forecasting*, **20**, 1048–1060.

Jensen, N., 1983: A note on wind generator interaction. Technical report, Risø-M-2411, Risø DTU.

Jónsson, T., P. Pinson, and H. Madsen, 2010: On the market impact of wind energy forecasts. *Energy Economics*, **32**, 313 – 320, doi:10.1016/j.eneco.2009.10.018.

Kankiewicz, A., J. Y. Li, and C. Finley, 2010: Turbine nacelle anemometers - a valuable resource for forecasting and operational assessment. *Poster presentation at AWEA Windpower*, Dallas, U.S.A.

Kelley, N., M. Shirazi, D. Jager, S. Wilde, J. Adams, M. Buhl, P. Sullivan, and E. Patton, 2004: Lamar Low-Level Jet Project Interim Report. Technical Report NREL/TP-500-34593, National Renewable Energy Laboratory; Golden, CO, U.S.A.

Krohn, S., P. Morthorst, and S. Awerbuch, 2009: The economics of wind energy. Technical report, European Wind Energy Association.

Landberg, L., 1997: The availability and variability of the European wind resource. *International Journal of Solar Energy*, **18**, 313–320, doi:10.1080/01425919708914326.

Larsen, G. C., H. A. Madsen, F. Bingöl, J. Mann, S. Ott, J. N. Sørensen, V. Okulov, N. Troldborg, M. Nielsen, K. Thomsen, T. Larsen, and R. Mikkelsen, 2007: Dynamic wake meandering modeling. Technical Report Risø-R-1607, Risø DTU.

Larsén, X. G., J. Mann, J. Berg, H. Göttel, and D. Jacob, 2010: Wind climate from the regional climate model REMO. *Wind Energy*, **13**, 279–296.

Lazarus, S. M., M. E. Splitt, M. D. Lueken, R. Ramachandran, X. Li, S. Movva, S. J. Graves, and B. T. Zavodsky, 2010: Evaluation of data reduction algorithms for real-time analysis. *Wea. Forecasting*, **25**, 837–851, doi:10.1175/2010WAF2222296.1.

Leidner, S. M., D. R. Stauffer, and N. L. Seaman, 2001: Improving short-term numerical weather prediction in the California coastal zone by dynamic initialization of the marine boundary layer. *Mon. Wea. Rev.*, **129**, 275–294.

Li, X. and Z. Pu, 2008: Sensitivity of numerical simulation of early rapid intensification of hurricane Emily (2005) to cloud microphysical and planetary boundary layer parameterizations. *Mon. Wea. Rev.*, **136**, 4819–4838.

Liu, Y., F. Chen, T. Warner, and J. Basara, 2006: Verification of a mesoscale data-assimilation and forecasting system for the Oklahoma city area during the Joint Urban 2003 field project. *J. Appl. Meteor. Climatol.*, **45**, 912–929, doi:10.1175/JAM2383.1.

Liu, Y., W. Cheng, Y. Liu, G. Wiener, R. Frehlich, W. Mahoney, T. Warner, J. Himelic, K. Parks, and S. Early, 2010: Impact of assimilating met-tower, turbine nacelle anemometer and other intensified wind farm observation systems on 0 – 12h wind energy prediction using the NCAR WRF-RTFDDA model. *10th EMS Annual Meeting and 8th ECAC*, Zürich, Switzerland.

Liu, Y., T. Warner, Y. Liu, C. Vincent, W. Wu, B. Mahoney, S. Swerdlin, K. Parks, and J. Boehnert, 2011: Simultaneous nested modeling from the synoptic scale to the LES scale for wind energy applications. *Journal of Wind Engineering and Industrial Aerodynamics*, **99**, 308 – 319, doi:10.1016/j.jweia.2011.01.013.

Liu, Y., T. T. Warner, J. F. Bowers, L. P. Carson, F. Chen, C. A. Clough, C. A. Davis, C. H. Egeland, S. F. Halvorson, T. W. Huck, L. Lachapelle, R. E. Malone, D. L. Rife, R.-S. Sheu, S. P. Swerdlin, and D. S. Weingarten, 2008: The operational mesogamma-scale analysis and forecast system of the U.S. army test and evaluation command. Part I: Overview of the modeling system, the forecast products, and how the products are used. *J. Appl. Meteor. Climatol.*, **47**, 1077–1092, doi:10.1175/2007JAMC1653.1.

Marquis, M., J. Wilczak, M. Ahlstrom, J. Sharp, A. Stern, J. C. Smith, and S. Calvert, 2011: Forecasting the wind to reach significant penetration levels of wind energy. *Bull. Amer. Meteor. Soc.*, **92**, 1159–1171, doi:10.1175/2011BAMS3033.1.

Mass, C., D. Ovens, K. Westrick, and B. A. Colle, 2002: Does increasing horizontal resolution produce more skillful forecasts. *Bulletin of the American Meteorological Society*, **83**, 407–430.

Mellor, G. L. and T. Yamada, 1982: Development of a turbulence closure model for geophysical fluid problems. *Rev. Geophys. and Space Phys.*, **20**, 851–875.

Meng, Z. and F. Zhang, 2011: Limited-area ensemble-based data assimilation. *Mon. Wea. Rev.*, **139**, 2025–2045, doi:10.1175/2011MWR3418.1.

Mitchell, H. L., P. L. Houtekamer, and G. Pellerin, 2002: Ensemble size, balance, and model-error representation in an Ensemble Kalman Filter. *Mon. Wea. Rev.*, **130**, 2791–2808.

Nakanishi, M. and H. Niino, 2006: An improved Mellor-Yamada Level-3 model: Its numerical stability and application to a regional prediction of advection fog. *Bound.-Layer Meteor.*, **119**, 397–407.

Nolan, D. S., J. A. Zhang, and D. P. Stern, 2009: Evaluation of planetary boundary layer parameterizations in tropical cyclones by comparison of in situ observations and high-resolution simulations of Hurricane Isabel (2003). Part I: Initialization, maximum winds, and the outer-core boundary layer. *Mon. Wea. Rev.*, **137**, 3651–3674, doi:10.1175/2009MWR2785.1.

Ochotta, T., C. Gebhardt, D. Saupe, and W. Wergen, 2005: Adaptive thinning of atmospheric observations in data assimilation with vector quantization and filtering methods. *Quarterly Journal of the Royal Meteorological Society*, **131**, 3427–3437, doi:10.1256/qj.05.94.

Otte, T. L., 2008: The impact of nudging in the meteorological model for retrospective air quality simulations. Part I: Evaluation against national observation networks. *J. Appl. Meteor. Climatol.*, **47**, 1853–1867, doi:10.1175/2007JAMC1790.1.

Otte, T. L., N. L. Seaman, and D. R. Stauffer, 2001: A heuristic study on the importance of anisotropic error distributions in data assimilation. *Mon. Wea. Rev.*, **129**, 766–783.

Pannekoucke, O., L. Berre, and G. Desroziers, 2008: Background-error correlation length-scale estimates and their sampling statistics. *Quarterly Journal of the Royal Meteorological Society*, **134**, 497–508, doi:10.1002/qj.212.

Peña, A., S. E. Gryning, and C. B. Hasager, 2010: Comparing mixing-length models of the diabatic wind profile over homogeneous terrain. *Theor. Appl. Climatol.*, **100**, 325–335.

Peña, A. and A. N. Hahmann, 2011: Atmospheric stability and turbulence fluxes at Horns Rev — An intercomparison of sonic, bulk and WRF model data. *Wind Energy*, doi:10.1002/we.500.

Peña, A. D., 2009: *Sensing the Wind Profile*. ISBN: 978-87-550-3709-0, Risø DTU.

Pedersen, T., 2004: Characterisation and classification of RISØ p2546 cup anemometer. Technical report, Technical Report Risø-R-1364(ed.2), Risø DTU.

Pedersen, T. F., J. Gottschall, J. R. Kristoffersen, and J.-A. Dahlberg, 2010: Yawing and performance of an offshore wind farm. *European Wind Energy Conference*, Warsaw, Poland.

Pielke, R. A. S., 2002: *Mesoscale Meteorological Modeling*, volume 78 of *International Geophysics Series*. Academic Press, 2nd edition, 676 pp.

Pleim, J., 2007: A combined local and nonlocal closure model for the atmospheric boundary layer. Part I: Model description and testing. *J. Appl. Meteorol. Climatol.*, **46**, 1383–1395.

Pleim, J. E. and R. Gilliam, 2009: An indirect data assimilation scheme for deep soil temperature in the Pleim-Xiu land surface model. *J. Appl. Meteor. Climatol.*, **48**, 1362–1376, doi:10.1175/2009JAMC2053.1.

Pleim, J. E. and A. Xiu, 2003: Development of a land surface model. Part II: Data assimilation. *J. Appl. Meteor.*, **42**, 1811–1822.

Rabier, F., 2005: Overview of global data assimilation developments in numerical weather-prediction centres. *Quarterly Journal of the Royal Meteorological Society*, **131**, 3215–3233, doi:10.1256/qj.05.129.

Reen, B. P., 2007: *Data assimilation strategies and land-surface heterogeneity effects in the planetary boundary layer*. Ph.D. thesis, The Pennsylvania State University.

Reen, B. P. and D. R. Stauffer, 2010: Data assimilation strategies in the planetary boundary layer. *Boundary-Layer-Meteorology*, **137**, 237–269, doi:10.1007/s10546-010-9528-6.

Ruggiero, F. H., K. D. Sashegyi, R. V. Madala, and S. Raman, 1996: The use of surface observations in four-dimensional data assimilation using a mesoscale model. *Mon. Wea. Rev.*, **124**, 1018–1033.

Seaman, N. L., 2000: Meteorological modeling for air-quality assessments. *Atmospheric Environment*, **34**, 2231 – 2259, doi:10.1016/S1352-2310(99)00466-5.

Shin, H. H. and S.-Y. Hong, 2011: Intercomparison of planetary boundary-layer parametrizations in the WRF model for a single day from CASES-99. *Bound.-Layer Meteor.*, **139**, 261–281.

Shin, H. H., S.-Y. Hong, and J. Dudhia, 2011: Impacts of the lowest model level height on the performance of planetary boundary layer parameterizations. *Mon. Wea. Rev.*, **140**, 664–682, doi:10.1175/MWR-D-11-00027.1.

Skamarock, W. C., J. B. Klemp, J. Dudhia, D. O. Gill, M. G. Barker, Duda, X.-Y. Huang, D.M., W. Wang, and J. G. Powers, 2008: A description of the Advanced Research WRF version 3. Technical report, NCAR Tech Notes-475+STR.

Smaïli, A. and C. Masson, 2004: On the rotor effects upon nacelle anemometry for wind turbines. *Wind Engineering*, **28**, 695–714, doi:10.1260/0309524043729958.

Sommer, A., 2002: Programme for measuring wind, wave and current at Horns Rev. Technical Report T011746.01.06, Eltra, Tech-wise Report.
URL http://130.226.56.153/rispubl/NEI/nei-dk-4851.pdf

Sørensen, P., N. A. Cutululis, A. Vigueras-Rodríguez, H. Madsen, P. Pinson, L. E. Jensen, J. Hjerrild, and M. Donovan, 2008: Modelling of power fluctuations from large offshore wind farms. *Wind Energy*, **11**, 29–43, doi:10.1002/we.246.

Stauffer, D. R. and N. L. Seaman, 1990: Use of four-dimensional data assimilation in a limited-area mesoscale model. Part I: Experiments with synoptic-scale data. *Mon. Wea. Rev.*, **118**, 1250–1277.

— 1993: Multiscale four-dimensional data assimilation. *Journal of Applied Meteorology*, **33**, 416 – 434.

Stauffer, D. R., N. L. Seaman, and F. S. Binkowski, 1991: Use of four-dimensional data assimilation in a limited-area mesoscale model. II - Effects of data assimilation within the planetary boundary layer. *Monthly Weather Review*, **119**, 734–754.

Storm, B. and S. Basu, 2010: The WRF model forecast-derived low-level wind shear climatology over the United States Great Plains. *Energies*, **3**, 258–276.

Storm, B., J. Dudhia, S. Basu, A. Swift, and I. Giammanco, 2008: Evaluation of the Weather Research and Forecasting model on forecasting low-level jets: Implications for wind energy. *Wind Energy*, **12**, 81–90.

Stull, R. B., 1988a: *An Introduction to Boundary Layer Meteorology*, volume 13 of *Atmospheric and Oceanographic Sciences Library*. Springer, 680 pp., ISSN 978-90-277-2769-5.

— 1988b: *An Introduction to Boundary Layer Meteorology*, Springer, volume 13 of *Atmospheric and Oceanographic Sciences Library*, chapter 9. 680, iSSN 978-90-277-2769-5.

Sukoriansky, S., B. Galperin, and V. Perov, 2006: A quasi-normal scale elimination model of turbulence and its application to stably stratified flows. *Nonlinear Process. Geophys.*, **13**, 9–22.

Talbot, C., E. Bou-Zeid, and J. A. Smith, 2010: Impact of land-surface variability on turbulent surface fluxes in multi-scale atmospheric simulations. *19th Symposium on Boundary Layers and Turbulence*, Keystone, CO, U.S.A.

Taylor, K. E., 2001: Summarizing multiple aspects of model performance in a single diagram. *J. Geophys. Res.*, **106**, 7183–7192, doi:10.1029/2000JD900719.

Torn, R. D. and G. J. Hakim, 2008: Performance characteristics of a pseudo-operational ensemble kalman filter. *Monthly Weather Review*, **136**, 3947–3963, doi:10.1175/2008MWR2443.1.

Troen, I. and L. Mahrt, 1986: A simple model of the atmospheric boundary layer; sensitivity to surface evaporation. *Bound.-Layer Meteor.*, **37**, 129–148.

Troen, I. and E. L. Petersen, 1989: European Wind Atlas. ISBN: 87-550-1482-8, Risø National Laboratory, Roskilde, Denmark.

Tuohy, A., P. Meibom, E. Denny, and M. O'Malley, 2009: Unit commitment for systems with significant wind penetration. *IEEE Transactions on Power Systems*, **24**, 592–601, doi:10.1109/TPWRS.2009.2016470.

Wagner, R., 2010: *Accounting for the speed shear in wind turbine power performance measurement*. Ph.D. thesis, Risø DTU.

Wang, W., C. Bruyère, M. Duda, J. Dudhia, D. Gill, H.-C. Lin, J. Michaelakes, S. Rizvi, and X. Zhang, 2010: *WRF-ARW Version 3 Modeling System User's Guide*. Mesoscale & Microscale Meteorology Division, National Center for Atmospheric Research, Boulder, USA.

Warner, T. T., 2011: *Numerical Weather and Climate Prediction*. Cambridge University Press.

Wilks, D. S., 2006: *Statistical methods in the atmospheric sciences*, volume 91 of *International Geophysics Series*. Academic Press, 2nd edition.

Zahle, F. and N. N. Sørensen, 2011: Characterization of the unsteady flow in the nacelle region of a modern wind turbine. *Wind Energy*, **14**, 271–283, doi:10.1002/we.418.

Zervos, A. and C. Kjaer, 2008: Pure Power - wind energy scenarios up to 2030. Technical report, European Wind Energy Association.
URL http://www.ewea.org/fileadmin/ewea_documents/documents/publications/reports/purepower.pdf

Zhang, D.-L. and W.-Z. Zheng, 2004: Diurnal cycles of surface winds and temperatures as simulated by five boundary layer parametrisations. *J. Appl. Meteorol.*, **43**, 157–169.

Zupanski, D., K. Paquin, R. Kelly, S. Nelson, M. Zupanski, I. Jankov, and P. Mallapragada, 2010: Applications of data assimilation methodologies in wind power forecasting. *EGU General Assembly*, Vienna, Austria.

Publications within this thesis

Journal papers:

C. Draxl, L. Delle Monache, F. Vandenberghe, Y. Liu, A. N. Hahmann, 2012: Thinning strategies for the assimilation of wind farm observations to improve wind energy predictions. *to be submitted to Wind Energy*

C. Draxl, A. N. Hahmann, A. Peña, G. Giebel, 2011: Evaluating winds and vertical wind shear from WRF model forecasts using seven PBL schemes. *Conditionally accepted in Wind Energy*

Project reports:

G. Giebel, R. Brownsword, G. Kariniotakis, M. Denhard, **C. Draxl**, 2011: The State-Of-The-Art in Short-Term Prediction of Wind Power : A Literature Overview, 2nd edition, ANEMOS.plus deliverable report.

Conference Proceedings:

C. Draxl, A. N. Hahmann, A. Peña, J. N. Nissen, G. Giebel, 2010: Validation of boundary-layer winds from WRF Mesoscale Forecasts with applications to wind energy forecasting. Talk and proceedings at the *19th Symposium on Boundary Layers and Turbulence*, Keystone, CO (US)

C. Draxl, A. N. Hahmann, A. Peña, J. N. Nissen, G. Giebel, 2010: Validation of boundary-layer winds from WRF mesoscale forecasts with applications to wind energy forecasting. Poster and proceedings at the *9th International Workshop on Large-Scale Integration of Wind Power into Power Systems as well as on Transmission Networks for Offshore Wind Power Farms*, Quebec, Canada

C. Draxl, C.L. Vincent, A.N. Hahmann, G. Giebel, P. Pinson, 2009: Variability forecasts for wind farms using high resolution initial conditions. *8th International Workshop on Large-Scale Integration of Wind Power into Power Systems as well as on Transmission Networks for Offshore Wind Farms*, Bremen, Germany

C. L. Vincent, **C. Draxl**, G. Giebel, P. Pinson, J. Jørgensen, C. Möhrlen, 2009: Spectral Verification of a Mesoscale Ensemble. *European Wind Energy Conference*, Marseille, France

C. Vincent, **C. Draxl**, G. Giebel, P. Pinson, J. Jørgensen, C. Möhrlen, 2009 : Spectral verification of a mesoscale ensemble. *European Geosciences Union*, Vienna, Austria

Talks:

23 Sep 2010: University of Salt Lake City, invited Talk: *A glance at Risø's reseach activities in wind energy: WRF for wind energy forecasts*

25 Feb 2009: International Energy Agency Workshop Wind Farms – Wake Effects and Power Fluctuations, Roskilde, Denmark: *New data assimilation techniques for short-term wind energy forecast models with a rapid update cycle*

1 – 4 Aug 2011: 14th Conference on Mesoscale Processes/15th Converence on Aviation, Range, and Aerospace Meteorology (ARAM), Los Angeles: *On the impact of the assimilation of nacelle winds and yaw angles with WRF-FDDA and WRF-DART for short-term wind energy predictions.*
http://ams.confex.com/ams/14Meso15ARAM/techprogram/paper_191185.htm

Poster presentations:

C. Draxl, L. Delle Monache, Y. Liu, G. Descombes, F. Vandenberghe, W. Y. Y. Cheng, A. N. Hahmann, G. Giebel, and J. C. Knievel, 2011: On the impact of the assimilation of nacelle winds and yaw angles with WRF-FDDA and WRF-DART for short-term wind energy predictions. *12th WRF User's workshop*, Boulder, USA

A. N. Hahmann, **C. Draxl**, A. Peña, J. Refslund, 2011: Simulating the vertical structure of the wind with the weather research and forecasting (WRF) model. *European Wind Energy Conference*, Brussels, Belgium

A. Hahmann, J. Badger, A. Peña, I. Karagali, X. Larsen, **C. Draxl**, M. Kelly, 2010: Using the WRF Model for Wind Energy Resource Assessment. *11th WRF User's Workshop*, Boulder, USA

C. Draxl, A. Hahmann, A. Peña, G. Giebel, 2010: Validation of boundary layer winds from WRF mesoscale forecasts with application to wind energy forecasting. *11th WRF User's Workshop*, Boulder, USA

A. N. Hahmann, C. L. Vincent, **C. Draxl**, 2009: A real-Time forecast system for Denmark: Verification of PBL processes for wind power forecasting and resource assessment. *Joint NCAR-NCAS WRF workshop*, Cambridge, UK

A. N. Hahmann, C. L. Vincent, **C. Draxl**, 2009: A real-Time forecast system for Denmark: Verification of PBL processes for wind power forecasting and resource assessment. *10th WRF Users' Workshop*, Boulder, USA