UPPSALA
UNIVERSITET

Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology 1428

Modelling Wind Power for Grid Integration Studies

JON OLAUSON

ACTA
UNIVERSITATIS
UPSALIENSIS
UPPSALA
2016

ISSN 1651-6214
ISBN 978-91-554-9690-6
urn:nbn:se:uu:diva-302837

Dissertation presented at Uppsala University to be publicly examined in Polhemsalen, Ångströmlaboratoriet, Lägerhyddsvägen 1, Uppsala, Friday, 4 November 2016 at 09:15 for the degree of Doctor of Philosophy. The examination will be conducted in English. Faculty examiner: Professor Thomas Hamacher (Technische Universität München).

Abstract
Olauson, J. 2016. Modelling Wind Power for Grid Integration Studies. *Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology* 1428. 114 pp. Uppsala: Acta Universitatis Upsaliensis. ISBN 978-91-554-9690-6.

When wind power and other intermittent renewable energy (IRE) sources begin to supply a significant part of the load, concerns are often raised about the inherent intermittency and unpredictability of these sources. In order to study the impact from higher IRE penetration levels on the power system, integration studies are regularly performed. The model package presented and evaluated in Papers I–IV provides a comprehensive methodology for simulating realistic time series of wind generation and forecasts for such studies. The most important conclusion from these papers is that models based on coarse meteorological datasets give very accurate results, especially in combination with statistical post-processing. Advantages with our approach include a physical coupling to the weather and wind farm characteristics, over 30 year long, 5-minute resolution time series, freely and globally available input data and computational times in the order of minutes. In this thesis, I make the argument that our approach is generally preferable to using purely statistical models or linear scaling of historical measurements.

In the variability studies in Papers V–VII, several IRE sources were considered. An important conclusion is that these sources and the load have very different variability characteristics in different frequency bands. Depending on the magnitudes and correlations of these fluctuation, different time scales will become more or less challenging to balance. With a suitable mix of renewables, there will be little or no increase in the needs for balancing on the seasonal and diurnal timescales, even for a fully renewable Nordic power system. Fluctuations with periods between a few days and a few months are dominant for wind power and net load fluctuations of this type will increase strongly for high penetrations of IRE, no matter how the sources are combined. According to our studies, higher capacity factors, more offshore wind power and overproduction/curtailment would be beneficial for the power system.

Keywords: Wind power, Wind power modelling, Intermittent renewables, Variability, Integration or renewables, Reanalysis data, Power system studies

Jon Olauson, Department of Engineering Sciences, Electricity, Box 534, Uppsala University, SE-75121 Uppsala, Sweden.

ISSN 1651-6214
ISBN 978-91-554-9690-6
urn:nbn:se:uu:diva-302837 (http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-302837)

Dedicated to the Royal Academic Orchestra

List of papers

This thesis is based on the following papers, which are referred to in the text by their Roman numerals.

I **J. Olauson** and M. Bergkvist, "Modelling the Swedish wind power production using MERRA reanalysis data", *Renewable Energy*, vol. 76, pp. 717–725, 2015

II **J. Olauson**, H. Bergström and M. Bergkvist, "Restoring the missing high-frequency fluctuations in a wind power model based on reanalysis data", *Renewable Energy*, vol. 96, pp. 784–791, 2016

III **J. Olauson**, M. Bergkvist and J. Rydén, "Simulating intra-hourly wind power fluctuations on a power system level", Resubmitted after revisions to *Wind Energy*, August 2016

IV **J. Olauson** and M. Bergkvist, "A new approach to obtain synthetic wind power forecasts for integration studies", Resubmitted after revisions to *Energies*, September 2016

V J. Widén, N. Carpman, V. Castellucci, D. Lingfors, **J. Olauson**, F. Remouit, M. Bergkvist, M. Grabbe and R. Waters, "Variability assessment and forecasting of renewables: A review for solar, wind, wave and tidal resources", *Renewable & Sustainable Energy Reviews*, vol. 44, pp. 356–375, 2015

VI **J. Olauson** and M. Bergkvist, "Correlation between wind power generation in the European countries", *Energy*, vol. 114, pp. 663–670, 2016

VII **J. Olauson**, N. Ayob, M. Bergkvist, N. Carpman, V. Castellucci, A. Goude, D. Lingfors, R. Waters and J. Widén, "Net load variability in the Nordic countries with a highly or fully renewable power system", Resubmitted after revisions to *Nature Energy*, August 2016

Reprints were made with permission from the publishers.

The author has also contributed to the following work, not included in the thesis.

A **J. Olauson**, A. Goude and M. Bergkvist, "Wind energy converters and photovoltaics for generation of electricity after natural disasters", *Geografiska Annaler: Series A, Physical Geography*, vol. 97, no. 1, pp. 9–23, 2015

B **J. Olauson**, J. Samuelsson, H. Bergström and M. Bergkvist, "Using the MIUU model for prediction of mean wind speed at low height", *Wind Engineering*, vol. 39, no. 5, pp. 507–518, 2015.

C **J. Olauson**, D. Lingfors, M. Bergkvist and J. Widén, "Quantifying variability: A review of metrics and a case study of net load variability", *Proceedings for 13th wind integration workshop*, Berlin, Germany, 2014.

D **J. Olauson**, H. Bergström and M. Bergkvist, "Scenarios and time series of future wind power production in Sweden", *Energiforsk report 2015:141*, ISBN 978-91-7673-141-3, 2015.

Contents

Acknowledgements

And since dissertations can be written about everything under the sun, the number of topics is infinite. Sheets of papers covered with words pile up in archives sadder than cemeteries [...]

–M. Kundera, *The unbearable lightness of being*, 1984

Performing research can sometimes be a lonely task and, as Kundera kindly reminds us, one can get the feeling that the output is often archived and forgotten. I however consider myself lucky in being recommended a research topic which I find highly interesting and rewarding and with possibilities for fruitful collaborations in and outside academia. For this, I am very thankful to my main supervisor Mikael Bergkvist. I would also like to thank Mikael for always finding time in his busy schedule for discussions and for contributing to the good ambience at our division.

My co-supervisor Mats Leijon is a busy man with many irons in the fire. Although Mats has not been involved in the details of my research, he is acknowledged for obtaining funding for my work and for the confidence he puts in us PhD students. MSB (Swedish civil contingencies agency) and Energiforsk are gratefully acknowledged for providing research funding. I hope you find it money well spent.

Research is much more fun if you don't do it yourself! I would like to thank my co-authors Anders, Mikael, Joakim W, Nicole, Valeria, David, Flore, Mårten, Rafael, Jonatan, Hans, Nasir, Jesper, Johan and Joakim L. I hope we will have the chance to work more together in the near future. Thanks also go to teaching colleagues and students, to the administrative personnel Ingrid, Maria, Gunnel, Thomas and Anna for always being helpful, to Colin for sharing his expertise on Belgian wind power, to Magnus and Tobias for their help with filter design, to my office roommates Yue, Maria and Tobias, to the organizers, senior researchers and PhD students at CNDS (Centre for Natural Disaster Sciences) and to the people in the wind power group: Hans, Sandra, Mikael, Anders, Per, Morgan, Stefan, Petter, Bahri, Marcus, Erik, Senad and Victor.

I do not forget my former colleagues at Sweco: you taught me a lot. Thank you Fredrik for being a very flexible manager and to Mattias, Olle and Pelle.

I would also like to thank Oskar Sämfors and Erik Böhlmark (Svenska Kraftnät), Daniel Gustafsson, Pierre-Julien Trombe, Roy Lilleberg, Linn Saarinen and Tobias Nylander (Vattenfall), Lennart Söder and Richard Scharff (KTH), Anton Steen (Svensk Vindenergi), Matthias Mohr, Jon Kjellin and Morgan Rossander (Uppsala University), Åsa Elmqvist (Energiforsk) and Johanna Laakso (Energimyndigheten) for valuable input to our work. The following persons have had a particularly important impact on this thesis frame: Dorothée, Anders, and Mikael who commented on the manuscript and Hannes and Valny who made the beautiful cover (photographer back cover: Hans Bernhoff).

Many persons have also helped with data access and interpretation. To mention a few: Marcus Flarup and Sverker Hellström at SMHI, Mattias Wondollek at Svensk Vindenergi, Birger Fält at Svenska Kraftnät, Kari Fougman at Sweco and Karin Salevid at Vattenfall. The hard work of many weather services around the world have led to extremely important (and often freely available) resources for research. Researchers at NASA and NOAA are acknowledged for producing such datasets, which have been used in all papers included in this thesis. Transmission system operators sharing their data are also gratefully acknowledged.

A special thanks goes to my fellow musicians in the Royal Academic Orchestra and its committed leader Stefan Karpe. If it wasn't for you my life would have been so much duller and I would probably never have returned to Uppsala for PhD studies. This thesis is therefore dedicated to you. Any orchestra would soon collapse without a strong viola section; great work Chizu, Helga, Amanda, Tom, Anna, Samuel, Mari, Thomas, Ester, Maria, Johanna, Johan, Kalle, Geraldine, Laura, Daniel and Alice. And thank you Julia, Vendela, Emil, Matteo, Malin, Oskar, Lena, Erik, Anja, Moa, Andrea, Diego, Gustav, Jonatan, Viktor, Anna, Klara, Linda, Viktor B, Pelle, Mattias, Tomas, Beatriz, Jan, Björn, Bo, Anna-Maria etc. for our chamber music experiences.

Isak, Colin, Kalle, Frida, Jonatan, Mi-ae and Emil, thank you for wonderful hiking, skiing, paddling and cycling trips. And to old friends whom I haven't seen as much as I would have liked; Berk, Tomas, Amanda, Jens, Marja, Olle, Åsa and Jonas, I hope you can make it to the dissertation party. Finally I would like to thank Dodo and my dear family Ingrid, Erland, Hannes, Viktor and Linn for your love and support.

Abbreviations

ACF	Autocorrelation function
ARMA	Autoregressive moving average
CART	Classification and regression tree
CF	Capacity factor
EOF	Empirical orthogonal function
FE	Forecast error
FFT	Fast Fourier transform
GB	Gradient boosting
GEFS	Global ensemble forecasting system
HVDC	High-voltage direct current
IEA	International Energy Agency
IFFT	Inverse fast Fourier transform
IRE	Intermittent renewable energy
LTC	Long-term correction
MA	Moving average
MERRA	Modern-era retrospective analysis for research and applications
MIUU	Department of meteorology, Uppsala University
NWP	Numerical weather prediction
PC	Power curve
PSD	Power spectral density
p.u.	Per unit
PV	Photovoltaics
RMS	Root mean square
SD	Standard deviation
SE	Sweden (SE1–4 for the four different bidding zones)
SMHI	Swedish Meteorological and Hydrological Institute
SNSP	System non-synchronous penetration
TSO	Transmission system operator
VI	Variability index
WT	Wind turbine

1. Introduction

If we combine the equations of Magnetic Force (B) with those of Electric Currents (C) [...]

–J.C. Maxwell, *A Dynamical Theory of the Electromagnetic Field*, 1865

In the early history of wind power, the wind turbines (WTs) were built to supply local energy demands including e.g. water pumping and grinding. When it became more common for WTs to be connected to large electric grids in the 1970's, a rapid increase in tower height, turbine diameter and generator rating begun, which has continued ever since. [1] The main drivers of this development have been the better wind conditions at higher hub heights and economics of scale, which give lower costs per produced energy unit for larger WTs. The dominating design of modern WTs is three-bladed horizontal axis, but other concepts are also under consideration. Although the lion's share of todays installed capacity is grid-connected and large-scale, there is some interest in small-scale wind power. In most areas of the world, photovoltaics (PV), for which generation costs are much less dependent on scale, however outperforms small-scale wind [2].

Starting from a global installed capacity of 5 GW at the end of 1995 (which has now been surpassed by Sweden alone), the capacity reached 59 GW in year 2005 and 433 GW in 2015[1]. Currently, roughly 4% of the world's electricity demand is met by generation from WTs. In some countries and regions, wind power supplies a more substantial part of the load. Two examples are Spain and Denmark with 19% and 42% wind penetration[2]. By including planned and projected installations, several other regions will have high shares of wind power in the near future. The Global Wind Energy Council (GWEC) has published three long-term scenarios for wind deployment [3]. In the "moderate" scenario, the capacity will supersede 700 GW by 2020 and 2000 GW by 2040. Whether these projections will be realised or not is of course impossible to tell. One can however note that the GWEC moderate scenario published in 2008 has so far been too conservative.

[1] `http://www.gwec.net/` Accessed: 2016-04-20.

[2] Data from the system operators (Red Eléctrica and Energinet respectively) for year 2015.

Table 1.1. *Levelised cost of electricity for selected generation technologies in Sweden 2014 (excluding policy instruments, 6% cost of capital). Results from [4] recalculated with an exchange rate of 9.3 SEK/€. The rated capacities are the sizes of the (set of) power plants for which costs have been calculated, e.g. a 150 MW onshore wind farm.*

Technology	Rated capacity (MW)	Investment (€/W)	Energy cost (€/MWh)
Onshore wind	150	1.3	55
Offshore wind	600	2.5	81
Rooftop PV	0.005	1.7	183
PV farm	1	1.1	100
Nuclear	1600	4.3	58
Hydro power	90	2.2	49
Gas co-gen	150	0.75	63

According to a recent (year 2014) estimate for Sweden [4], onshore wind power compares favourably to most other generation technologies in terms of levelised cost of electricity. Table 1.1 shows a selection of the investment and energy costs, assuming 6% cost of capital, an exchange rate of 9.3 SEK/Euro and excluding policy instruments (taxes and subsidies). The costs include those for capital, operation and maintenance and insurances but not integration costs induced by intermittency and unpredictability. If policy instruments are taken into consideration, the competitiveness of wind power is increased further.

Wind power in Sweden and the Nordic synchronous power system are the main concerns of this thesis. Fig. 1.1 shows the eleven bidding zones[3] in the Nordic system and wind farms in operation or under construction as of April 2015. The data for Sweden are relatively accurate, but of somewhat poorer quality for the other countries. The database, taken from Paper VII, comprises 800–1000 MW in eastern Denmark[4], Finland and Norway respectively and 6500 MW in Sweden. As can be seen in the figure, the farms in Denmark and Sweden are well dispersed while deployment in Finland and Norway has been concentrated to the coastal areas. Recently (early 2016), an investment decision was taken for a 1000 MW project in bidding zone NO3. The six farms in this project, not included in the map, are to be completed by 2020. Maximum net transfer capacities[5] in MW for import to Sweden are given in red text in Fig. 1.1 and Swedish HVDC links are indicated by black arrows.

[3]Because of bottlenecks in the transmission system, the Nordic system has been divided into bidding zones. By allowing the electricity price to differ between the zones, actors are encouraged to act accordingly, e.g. install more capacity in zones with deficits. The current eleven zones have been in place since 2011.

[4]Western Denmark is not part of the Nordic synchronous system.

[5]`http://nordpoolspot.com/globalassets/download-center/tso/max-ntc.pdf` Accessed: 2016-04-25.

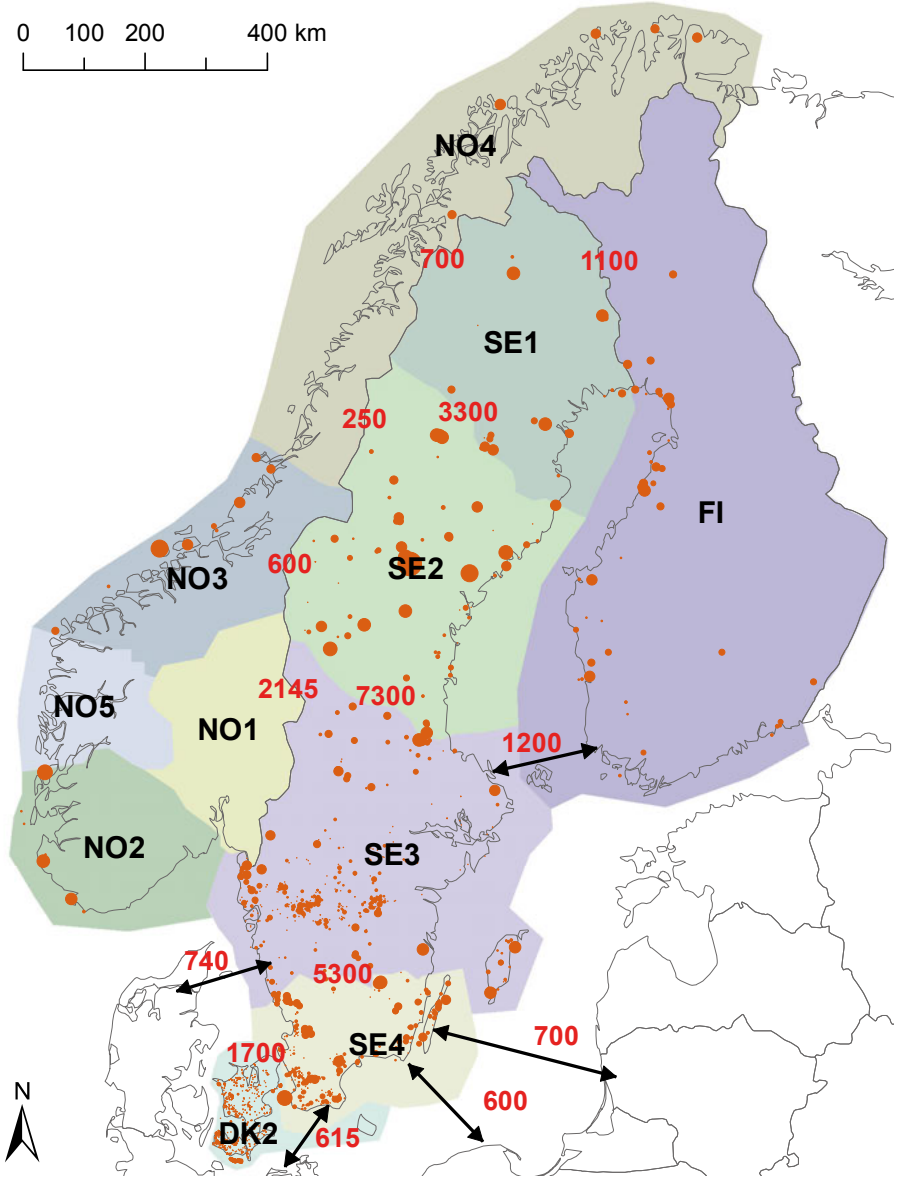

Figure 1.1. The eleven bidding zones and wind farms in the Nordic synchronous system (operating or under construction as of April 2015). The circle areas are proportional to farm capacities. Offshore country borders are only illustrative. Maximum net transfer capacities in MW for Sweden as of March 2016 are given in red text and Swedish HVDC links are indicated by black arrows. For comparison, the average load in Sweden is around 15,500 MW.

1.1 Wind integration challenges

Due to its variable and unpredictable nature, integration of wind power into the power system poses several challenges and induces additional costs [5]. Depending on the characteristics of the power system, e.g. thermal or hydro dominated, system size, load characteristics and transmission capacities, the magnitudes of the challenges for similar wind penetration levels differ.

By definition, there is always a balance between generation and load. If the load is high and the wind generation is low, sufficient capacity must be available from other generators, import and demand side flexibility. Assuming that conventional power plants are dismantled when wind farms are constructed, an important issue is therefore to make sure that the load can be supplied at all times. The deployment of wind power can also increase the magnitudes of the changes in net load[6], which may be an issue. Very short wind power fluctuations due to turbulent gusts are close to uncorrelated between sites and thus effectively dampened by aggregation. The ramps associated with passages of large-scale weather systems are most critical according to [6]. In addition to the inherent variability, wind power generation is difficult to predict, especially for longer forecast horizons. The planning of the power system thus becomes more challenging if the wind penetration is high.

The variability of wind power is not only a technical question, it also profoundly affects the energy markets and the profitability of both conventional generators and wind farms. By introducing intermittent generation with very low operating costs, the electricity price is reduced but also becomes more volatile. Systems with relatively high shares of intermittent generation, e.g. Germany and Denmark, sometimes even experience periods with negative prices. For operators of conventional plants, the introduction of more wind power may lead to lower capacity factors (CFs)[7] and more cycling. WT owners, on the other hand, may find that the average revenue for their sold electricity is lower than the time-averaged price [7]. A different generation portfolio implies that different market structures may be necessary. As an example, capacity markets are already in place in some systems and under consideration in others [8].

For several reasons, wind power requires investments in new and existing transmission grids. Firstly, wind farms are often located in remote areas, and thus requires construction of new lines. At least in Sweden, this expense is covered by the farm owner, so it should not be considered an integration cost. Investments in the existing grid can be necessary since wind farms have lower CFs than conventional generation (so larger transmission capacity is needed to transfer the same annual energy) and may be located far from load centres. Again with Sweden as example, load is concentrated in the southern part of the country, but wind farms are more evenly spread (see Fig. 1.1). Grid rein-

[6]The net load is generally defined as load minus intermittent renewable generation.

[7]The capacity factor is the mean generation normalised to the installed capacity.

forcements will become necessary in order to transfer wind power from north to south. In order to cope with the increased net load variability, it may also be rational to strengthen cross-border transmission capacity [9, 10]. Lastly, if conventional generation is dismantled when more WTs are erected, there may be deficits of reactive power in some regions. This leads to a less stable voltage and reduces the transmission capacity of existing lines.

Last but not least, replacing synchronous generators with wind farms decoupled from the grid by power electronics reduces the inertia of the system. The lower the inertia, the less damping is provided and the frequency thus reacts faster to abrupt changes in load or generation. From the system operators' point of view, the reduction of inertia is generally considered the main challenge with integration of renewables [11]. The instantaneous wind penetration levels (wind/load) can be considerably higher than the average. As an example, in Sweden 2015 the mean penetration was 12% but the maximum hourly penetration was 39%. In order to keep the inertia at acceptable levels, limitations can be set on the instantaneous penetration. From a technical point of view, it is not complicated to curtail wind power. If the average penetration is not very high, relatively small amounts of energy have to be spilled at times when prices can be expected to be low, see Section 4.2.2.

On the positive side, the dimensioning errors can be smaller for a system without large centralised generators such as coal or nuclear. Wind power can also enhance voltage and transient stability by providing reactive power [12], have fault-ride-through capabilities [13] and provide synthetic inertia [11]. It is likely that in a future with a significantly higher wind penetration, the farms will (be required to) have these capabilities and assist in balancing the power system. More ideas on how to tackle challenges with wind integration can be found in two reports [14, 15] in Swedish.

1.2 Research objectives and previous work

Over the last decades, the levelized cost of electricity (LCOE) for wind power has fallen steadily. According to [16], the global weighted LCOE for onshore wind was around 0.07 $/kWh year 2015, which can be compared to 0.20 $/kWh year 1995 (assuming a capital cost of 7.5%). For utility-scale PV, the decline has been even faster and the cost is now around 0.12 $/kWh. For year 2025, global weighted costs for onshore wind and PV are projected to be 0.05–0.06 $/kWh. It is thus reasonable to believe that integration challenges rather than LCOEs will be the main barriers to overcome if intermittent renewables should deserve a significant role in future power systems.

In order to achieve a smooth and effective transition to highly or fully renewable power systems, integration studies are necessary. Realistic wind power scenarios and corresponding time series, both generation and forecasts, are key components in any such study. The main research objective of this the-

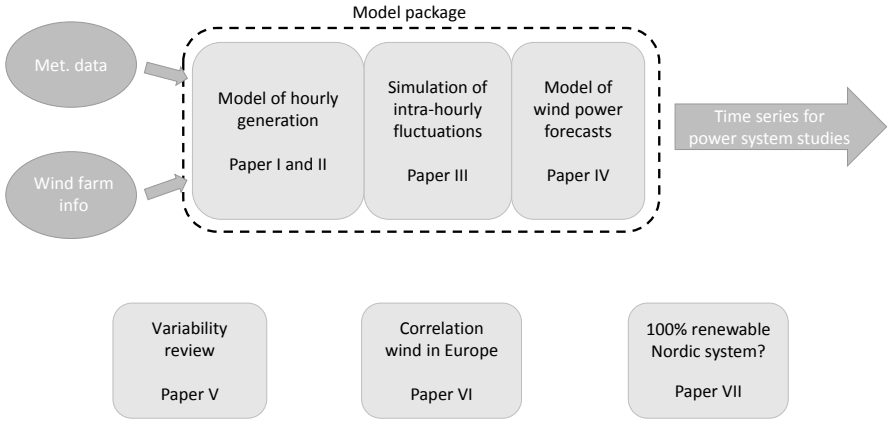

Figure 1.2. Scope of thesis. "Met. data" is short for meteorological data. Wind farm information can be either on actual farms (when the objective is to evaluate the models) or scenarios of future farms.

sis is to develop and improve methodologies for producing these time series. Contrary to e.g. [17], which uses highly resolved meteorological models, we rely on globally and freely available data from coarse, meteorological reanalyses[8]. This approach has its obvious benefits, but also requires that particular care is taken to ensure that e.g. high-frequency fluctuations are appropriately represented. A secondary research objective is the study of variability of intermittent renewables, possibly in combination with the electric load, and possibilities for reducing this variability. It is important to stress that these studies are statistical in nature and do not take into account transmission limitations, costs, balancing capabilities of hydro power etc. Such issues will be the topic for future work, see Chapter 6.

The scope of the papers included in this thesis is illustrated in Fig. 1.2. Methods and results from these papers are presented in this thesis frame. Additionally, selected results on wind power in Sweden are given. These are partially new and partially taken from Paper D.

In the remainder of this subsection, some work that has been influential for my thinking are credited. First, IEA Wind Task 25 [5, 18–20] on "Power systems with large amounts of wind power", led by Hannele Holttinen, should be acknowledged for their work on compiling results from wind integration studies and providing guidelines for such studies including input data requirements.

Freely available meteorological reanalyses [21, 22] have become popular for creating wind power time series [10, 23–28]. The complexity of the methods used and the extent of validations with measurements differ between the

[8]Long meteorological time series produced with a consistent model and assimilation system, see Section 2.2.

authors. Output from such models can be post-processed for obtaining appropriate high-frequency fluctuations or time series with higher temporal resolution [17, 29, 30]. In many types of wind integration studies, simulated (synthetic) forecasts are necessary for determining the impact from wind power uncertainty on the operation of power systems [18]. Our methodology for generating such time series differed from earlier attempts, but the multivariate ARMA method used by Söder [31] and others influenced our thinking.

Regarding statistical techniques for modelling and studying the variability of wind power, power spectral density estimates [32, 33], machine learning [34] and empirical orthogonal functions [35] have all been used earlier. Correlations of outputs between different farms or regions, which are important for wind power smoothing, have been studied by e.g. [36–38].

Finally, the study by Staffell and Green [23] on deterioration of wind farm performance and the work by Hirth *et. al.* [7,39] on the diminished revenues of wind power for higher penetration rates are acknowledged for their originality and high quality.

1.3 Guiding principles

In hindsight, a few guiding principles can be distinguished behind all papers included in the thesis. In conclusion:

- Focus on wind power on the power system level.
- The methods and results should be directly useful (engineering approach) and results should be shared.
- Use physical models as far as possible. When needed, statistical models are subsequently employed.
- Always validate the models with measurements on the power system level.
- Use as long time series as possible.
- Utilise, if possible, input and validation data that are publically and globally available.

Although large farms can have an effect on the regional grid [40], the largest impact from wind power is on the power system scale. This scale has thus been the focus of all papers included in the thesis. Better results are generally obtained if generation and forecasts are first modelled for each farm. When tuning the parameters for the models, measurements on the power system level were however the most important training data. By using freely available input data, demonstrating the methodology on a power system level (including reasonable computational times) and sharing the results, the chance that the methodologies are actually used by practitioners is increased.

A very important question is whether to use linear scaling of historical measurements [6,41,42], physical models or statistical models [43,44] for obtain-

ing data for power system studies. Our take on this matter is to use physical models as far as possible, employ statistical models for fine-tuning and only use measurements for validating the results. Purely statistical models were rejected since it is often crucial to have synchronous load and wind power data [18, 45]. Moreover, physical models are more credible for modelling future WTs and farms in new geographical areas (e.g. offshore).

Nobody questions that (quality controlled) measurements are valuable for studying the past. As the following example shows, linearly scaled measurements can however be very misleading for studying the future. For illustration, a paper by Kiviluoma *et. al.* [6] is used. The paper was chosen partly because it contains a readable analysis of wind power variability in several different countries, but also because the authors have some influence in the wind integration community and express doubts whether modelling is appropriate. In Section 3.2 in [6], net load ramps in future power systems with high wind penetrations were studied. Historical time series were linearly scaled since *"geographical smoothing is already incorporated in the data"*. According to our analyses in Paper D, this is to some degree correct; if more wind farms of the same type as already present are deployed in a mature power system, the reduction in normalised variability is negligible. However, the WTs have evolved over the years and can be expected to evolve in the future. This will have a large impact on the variability. As an example (see Section 4.2.2 for more results), let us consider the first percentiles of one hour step changes[9] in the Swedish net load. With no wind power, the first percentile for the years 2007–2014 was -1.14 GW. By adding 50 TWh/a of wind power, linearly scaled from historical time series, this figure is reduced to -1.58 GW. With a time series from a scenario of future deployment (scenario C1 from Paper D), -1.34 GW results. The increase in extreme negative ramp magnitudes is thus less than half of that obtained by using scaled measurements.[10]

Since the question of modelling versus measurements is so important, some results from modelling Swedish wind generation are given already now. As can be seen in Fig. 1.3, a physical model based on coarse reanalysis data and some statistical fine-tuning can adequately predict the actual generation. As for the whole thesis, different time periods were used for training the model and evaluating the performance.

As can be anticipated from the discussion above, the normalised variability of future wind power will most likely be significantly lower than today's. When studying the variability and related costs for wind power in the future,

[9]The change in averaged power over a fixed number of time steps, e.g. change in hourly averaged power over four hours.

[10]For the sceptical reader it can be worth pointing out that modelled time series based on historical farms yields in principle identical results as measurements (around 1% difference). The difference is thus not due to modelling per se, but rather due to the expected increase in CFs for future WTs. Scaling historical measurements implicitly assumes no change in WT characteristics.

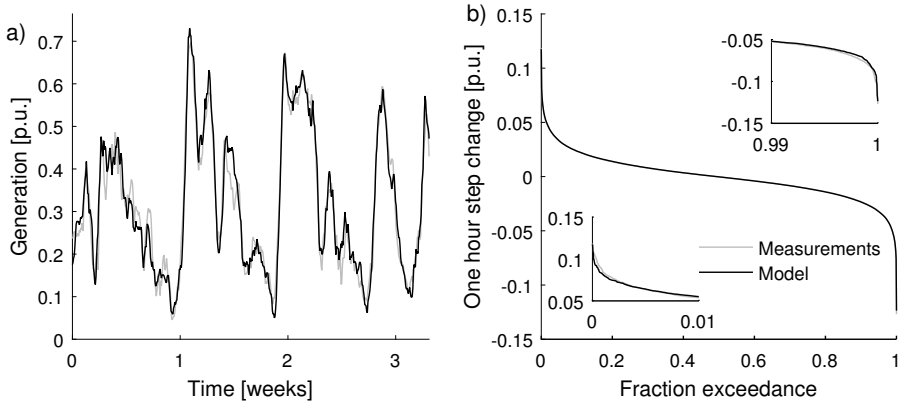

Figure 1.3. Models can successfully predict wind generation: **(a)** time series example, **(b)** step change duration curves with extremes as insets. Based on measurements and model output for the whole of Sweden (results from Paper II).

especially if decisions are to be taken about desirable development paths, I believe that these must be compared to the energy production, not the installed capacity. The former is much more informative regarding the benefits from wind power. Ten GW of wind power with a high CF will be somewhat more variable than ten GW of wind power with a low CF. This must not lead us to the false conclusion that it would be better to install low CF farms; if we want 30 TWh/a of electricity, the variability becomes significantly lower if high CF farms are deployed.

This section is concluded by a few remarks on hardware and software. A standard desktop computer from year 2011 was used for all computations. Since some of the datasets were quite large, an extra hard drive and 16 GB RAM memory have been installed. Matlab has been used solely for the modelling, except in Papers B and III were R was utilised to some extent. The latter program is free of charge and the preferred choice by most statisticians, so new methods generally appears first in R. Matlab, on the other hand, is faster, more user friendly, better documented and has nice graphical qualities (especially with the newly introduced colour schemes). No 3D-plots have been used since these, in my opinion, are most often only confusing.

1.4 Nomenclature

As much as I would have liked it, the nomenclatures in the different papers are not entirely consistent. Sometimes similar techniques have been used in several papers, but with different variable notations. In the method section of

this thesis frame, the nomenclature might thus be different from that in some of the papers.

The original purpose of my PhD project was to design and construct a small vertical axis wind turbine, see Section 1.6. In order to avoid confusion of the rotor of the generator, the "aerodynamic rotor" (rotor blades and hub) and the whole system (commonly denoted wind turbine), it was decided to denote the latter wind energy converter (WEC). This denomination was used in the earlier papers and in my licentiate thesis [2]. In the later papers and in this thesis, wind turbine (WT) is used.

As for the energy generated by WTs, different terminologies and units have been used. Average generation for certain fixed time periods has been denoted "hourly energy", "generation" or "(average) power" and units have been per unit (p.u.)[11], Watt etc. If the term power was chosen, this should always be interpreted as average power at sampling frequency.

When discussing a "wind power scenario", this refers to a scenario of both technical parameters of future WTs (capacity factor, specific rating[12] etc.) and the geographical distribution, i.e. coordinates of the farms. The term "aggregated" is used exclusively for geographical aggregation, e.g. the combined wind power generation in a country. Statistical metrics such as standard deviation and correlation are typically denoted by the symbol for the expected value instead of the sample metric, i.e. σ and ρ instead of s and r. In the thesis frame, variables and parameters are defined in their context. Abbreviations are listed on page 11.

1.5 Wind power research at UU

At the Division of Electricity at Uppsala University, research on straight-bladed, vertical axis WTs has been conducted for over ten years, focusing on e.g. generators [46, 47], aerodynamics [48], noise [49] and the electrical system [50]. Three WTs have been designed and constructed for research purposes:

- Torsholm 200 kW [51] - Constructed by the spin-off company Vertical Wind, but now owned by Uppsala University. A picture of Torsholm is shown on the back cover of this thesis.
- Lucia 12 kW [52] - Research wind turbine.
- Birgit 10 kW [53] - A novel concept turbine developed to electrify telecommunication towers.

Wind power studies are also conducted by meteorologists at the department of earth sciences. Examples of research fields are the MIUU model (see

[11] One per unit corresponds to the maximum theoretical generation, e.g. if the installed capacity is 1000 MW and 300 MWh is generated a certain hour, this equals 0.3 p.u.

[12] Rated power over swept rotor area.

Section 2.2.2), wind power in forests [54] and icing [55]. Paper II has a co-author from the department of earth sciences. In 2013, Gotland University was merged with Uppsala University. Parts of the island of Gotland have excellent wind conditions. Many of the first WTs in Sweden were deployed here, most notably the 3 MW, two-bladed "Matilda" in 1993. At Campus Gotland, the focus of technical wind power research is on wake modelling, see e.g. [56].

Wind power researchers at (primarily) Uppsala University and the Royal Institute of Technology (KTH) in Stockholm are organized in the network "STandUP for Wind". This centre, involving nearly 50 persons, is profiled towards planning and establishing wind energy in the Swedish electrical network.

1.6 Wind power and natural disasters

I am a member of the research school Centre for Natural Disaster Sciences (CNDS)[13]. My original research objective was to develop and construct a portable WT, intended for generation of electricity in the immediate response and for more long-term recovery work after a natural disaster. Based on the results from [2] and Papers A and B, it was concluded that photovoltaics is most often a better solution; the energy cost is generally lower and it is not so common with extended periods of very low generation, thus a smaller battery bank is needed. Exceptions are regions close to the poles, but these are not the most prone to natural disasters and are sparsely inhabited.

The abovementioned results lead to the new research direction that is presented in this thesis. Although the link to natural disasters may seem weaker than before, this is not the case. The whole electrical system is susceptible to disasters such as storms, earthquakes, extreme icing events etc. One of the eight key risks with a warmer climate pointed out by the IPCC [57] is *systemic risks due to extreme weather events leading to breakdown of infrastructure networks and critical services such as electricity, water supply and health and emergency services*. Moreover, the vulnerability of the power system to natural disasters can be altered by the power mix; introducing large amounts of highly variable production with little or no inertia jeopardise the stability of the grid if proper measures are not taken. In conclusion, wind power and natural disasters are linked in both directions and proper modelling of future wind power will be crucial for securing a robust power system.

[13] `http://www.cnds.se` Accessed: 2016-08-04.

2. Data

In this chapter, the data used in our studies are presented and some pitfalls of using coarse reanalyses for modelling wind power are discussed. Some datasets that were ultimately *not* utilized are also briefly mentioned. The data fall into two main categories: measurements and output from meteorological models.

2.1 Measurements

One of the guiding principles stated in Section 1.3 is that model outputs must always be compared to actual measurements. If the aim is to generate wind power time series for a power system, measurements for the whole system should primarily be used. In some cases it can however be advantageous to also utilise measurements for single farms, e.g. for determining correlations of forecast errors.

Time series of hourly, aggregated wind generation for the four bidding zones in Sweden are available at the Swedish transmission system operators (TSO) web page [58]. The data were normalised to the installed capacity at each time step, i.e. expressed in p.u. The definition of installed capacity is not clear-cut. Swedish Wind Energy[1] [59], includes a wind farm in their statistics when it has the "possibility to deliver power to the grid". The Swedish Energy Agency [60], on the other hand, only includes turbines that has been reported to the electricity certificate system. The two dates can be the same, but the latter is often up to a few months later since the owner does not want to report the turbine until the testing/trimming phase is completed. When the installed capacity increases significantly each year, the two databases can thus yield differences in the national capacity of a few percent. In Paper I, a considerable amount of work was put into correcting and supplementing the data from [60], thus producing a database of WTs installed year 2012 or earlier. For the analyses in Section 4.2, this database was extended to include farms deployed up until 2014. For the reason just mentioned, the installed capacity in this database may be an underestimation. An average of this time series and linearly interpolated end-of-year values from [59] was thus used[2].

[1] Svensk Vindenergi; the trade association for companies working with wind power in Sweden.
[2] Two other databases of wind farms are Vindstat and Vindbrukskollen, see `http://vindstat.com/` and `http://www.vindlov.se/sv/vindbrukskollen1/karta/`.

In some of the papers, data from other power systems were needed in addition to the Swedish. Time series of wind power generation, with temporal resolution ranging from five minutes to one hour, were thus retrieved from the system operators' web pages/servers or by mail. As for Sweden, the data were transformed to p.u. by normalising the MW values to the installed capacity at each time step. The resolution of the time series of installed capacity differed between the systems; for some, values for each day were available, but in some cases interpolation had to be used since only monthly or yearly values were available. All time series were quality-controlled, removing obviously erroneous data samples and, in some instances, filling the resulting gaps with interpolated data. When needed, the time series were shifted in time to get synchronous data. Most often, all data was converted to UTC+00 (coordinated universal time) or to UTC+01 (Swedish winter time). Particular care had to be taken to daylight saving time, which is handled differently in different databases.

Measurements from a few met masts and farms were obtained from the power company Vattenfall. Forecasted generation time series were shared by the same company. Data on electric load in the Nordic countries were attained from the respective TSO's web page. The Swedish load dataset, used in Paper VII and Section 4.2, is not entirely complete; electricity that is produced and consumed locally (e.g. in pulp industries) is not included[3].

2.2 Meteorological models

Numerical Weather Prediction (NWP) models are used to forecast wind speed and other variables for the present ("nowcasts") up to several days ahead. The same type of models can also be run with historical data from a constant data assimilation system in order to produce long and relatively consistent[4] datasets for research and climate services. Several of these retrospective analyses (reanalyses) are publicly available, e.g. ERA Interim [22], MERRA [21, 61] and JRA-55 [62]. See [63–65] for reviews that also include earlier generation reanalyses. This type of datasets has become popular for modelling wind power [10, 23–28] and was used in all papers included in this thesis. In Papers V and VII, solar irradiation and wave data from NWP reanalyses were also used. The following three subsections on coarse reanalyses, mesoscale models and forecasts respectively are however only devoted to wind power related variables.

[3]The lacking energy amounts to around 4% of the total load.

[4]With consistent we mean that the quality of the analysis/forecast does not change substantially over the years.

2.2.1 Coarse reanalyses

Because of the relatively high correlations with measurements [63] and temporal resolution, MERRA was the primary reanalysis choice in the papers included in this thesis. In Paper II, some variables were also taken from the ERA Interim dataset. CFSR data [66] were considered but ruled out since the time series are not consistent in time; from 2011 and onwards, only outputs from the operational model CFSv2 are available.

MERRA, Modern-Era Retrospective analysis for Research and Applications, is a NASA reanalysis for the satellite era, i.e. since 1979 [21]. Output is available with a spatial resolution of 0.5 degrees in latitude and 0.67 degrees in longitude and one hour temporal resolution[5]. The spatial resolution of MERRA is quite typical for modern global reanalyses. Available variables include e.g. northward and eastward wind speeds, temperature, humidity and air pressure. Wind speeds are available at two and ten meters above the displacement height (see Section 3.1) and at 50 m above ground. Recently, the atmospheric general circulation model was updated, resulting in a second generation reanalysis: MERRA2 [61]. The production of MERRA thus stopped in February 2016. MERRA2 has roughly the same resolution and spans the same time period as its predecessor. Results from citestaffell2016 show that the performances are similar in terms of modelling wind power generation.

Although a coarse reanalysis can be very useful for modelling the aggregated generation for a region or country, there are analyses that such a model should not be used for. A few examples are given in the numbered list below. For readers not familiar with wind power computations, it is recommended to read Sections 3.1 and 3.2 before proceeding.

1. Determine likely farm locations. As can be seen in Fig. 2.1, local (and some regional) variations in the mean wind speed are not captured by a coarse reanalysis. In [67], the geographical distribution of farms in Germany and the Nordic countries was optimised in terms of a low combined variability. Only sites with an estimated CF of 0.15 or higher were considered. The NWP data (11 km resolution) thus ruled out most regions in Sweden except an area in the south, the mountain region bordering Norway and some coastal sites. When comparing to the actual deployment and planned farms (see Fig. 4.15 on page 83), this picture is drastically wrong[6].
2. Calculate generation time series without first properly scale the wind speed. The actual mean wind speed at a site has little in common with

[5] Some variables, not used in our work, have a coarser temporal resolution.

[6] Note that [67] was chosen as an example since details on farms in Sweden were available for comparison and that the difference between prediction and actual outcome is so striking. Other authors [39, 68] have used reanalyses with an even coarser resolution for determining potential sites for deployment and/or compute generation time series without first properly scale the wind speed.

that obtained from coarse reanalysis datasets, in particular in more complex terrain. Because of the shape of the power curve, an error of 1 m/s in the mean wind speed gives around 20–40% error in the annual energy yield. The mean wind speed should thus be determined either from more highly resolved resource maps or from the anticipated annual energy yield of planned farms. Without a proper scaling, the variability in generation is also misjudged. If the mean wind speed is underestimated, the variability will be overestimated, see Fig. 2.2a for an example.

3. It is often desirable to perform a LTC (long-term correction) [69] of a shorter measurement, e.g. for estimating the mean output of a future farm or, as in [23], for quantifying deterioration in wind farm performance. In the latter case, the LTC is done in order to separate possible long-term trends in the wind climate from trends in farm performance. If a trend in the wind climate exists and an erroneous mean wind speed (e.g. taken directly from the coarse reanalysis) is assumed, an artificial trend in the farm performance will result, see Fig. 2.2b. Even if no long-term trend is present, the LTC will be more uncertain for shorter measurements (noisy signals in Fig. 2.2b).

2.2.2 Mesoscale models

In order to get more refined results, the output from global, coarse NWPs can be fed into regional models with higher resolution, a procedure referred to as dynamical downscaling. Typically, these mesoscale models have a spatial resolution ranging from a few hundred meters to 10 km. Results from a mesoscale model [70–72] developed at the former Department of Meteorology, Uppsala University (the MIUU model) were used in Paper B and for generation of scenarios of wind farm deployment in Paper D. Freely available data[7] on mean wind speeds can be retrieved for different heights and with spatial resolution 0.5 km, see Fig. 2.1a.

The Weather Research and Forecasting (WRF) model [73] is a NWP system intended for research and operational forecasts and is used by both major weather services and many universities. In the beginning of my PhD studies, we considered using WRF for producing more highly resolved data for modelling of wind power. For various reasons this never happened; firstly there is a significant learning threshold, secondly, as described in Section 1.3, we wanted as many as possible to be able to use our methodologies and thirdly, the fine results obtained (on the power system level) by using coarser models.

The Swedish Meteorological and Hydrological Institute (SMHI) uses mesoscale models for regional forecasts of various variables. Recently, a large amount of their data became publically available[8]. The MESAN(HIRLAM)

[7] `http://www.s1106835.crystone.net/?q=en/node/55` Accessed: 2016-07-15.
[8] `http://opendata-catalog.smhi.se/explore/` Accessed: 2016-04-14.

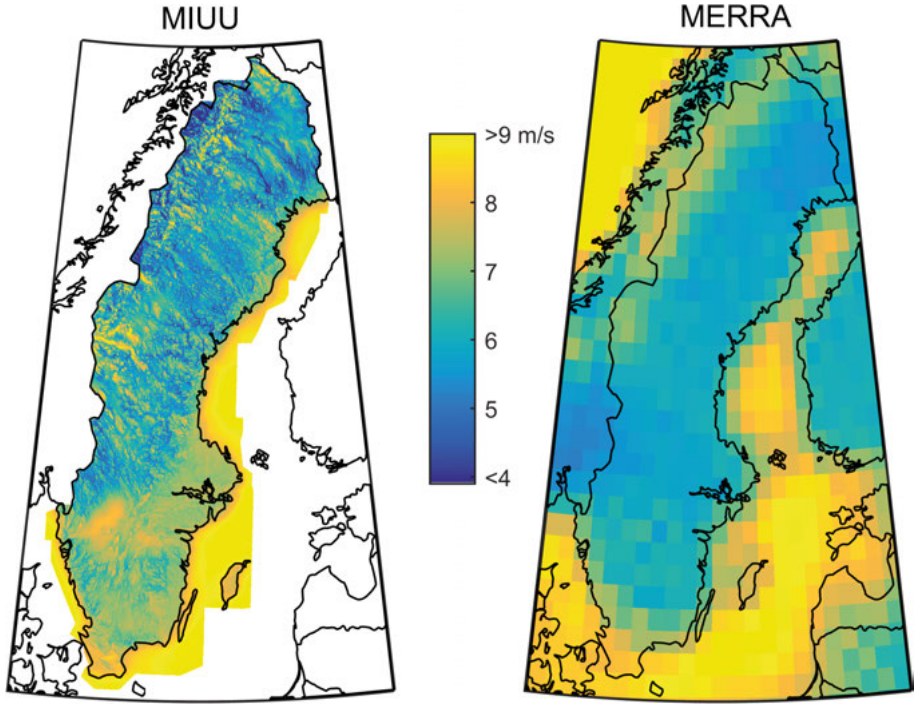

Figure 2.1. Mean wind speeds 100 m above displacement height according to the MIUU and MERRA models (see Sections 2.2.2 and 2.2.1 respectively). Horizontal resolutions are 500 m and $0.5 \times 0.67^{\circ}$ respectively. Local, and even some regional, variations are not captured by MERRA.

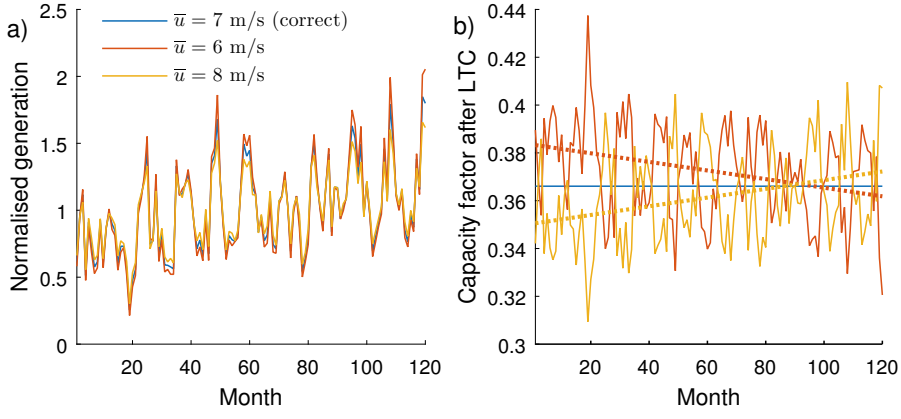

Figure 2.2. Example based on a MERRA time series for illustrating the effects of assumed mean wind speed on variability and long-term corrected generation. For a clear demonstration of the latter, a fictive linear trend was added to the time series, which was subsequently linearly scaled to means of 6, 7 and 8 m/s respectively. In this example it is assumed that 7 m/s is the correct mean wind speed. **(a)** The variability in normalised monthly generation is over- and underestimated with mean wind speeds of 6 and 8 m/s respectively. **(b)** When the monthly outputs based on a mean wind speed of 7 m/s are long-term corrected using time series with means of 6 and 8 m/s, artificial trends in LTC output results (dotted lines).

is an analysis model, describing the current state of the weather every hour in an eleven km grid. The MESAN dataset was found inconsistent, which is not surprising since, in contrast to reanalyses, the model has evolved over the years. As an example, the spatial resolution was not the same throughout the period. Furthermore, data were only available since 1998, i.e. significantly shorter than for MERRA. It was thus decided to not use MESAN-data.

2.2.3 (Re)forecasts

For creating synthetic wind power forecast errors with the methodology described in Section 3.2.5, forecasted wind speeds and other variables are necessary. Three sources of forecasts were considered and are described briefly below: regional forecasts produced by SMHI, global forecasts available in the TIGGE project and GEFS reforecasts. For readers not familiar with wind power forecasts, it may be useful to read Section 3.2.5 before proceeding.

SMHI uses a limited-area model to generate regional forecasts covering essentially the Nordic countries. SMHI models have a spatial resolution of between 2.8 and 11 km and forecast horizons up to 240 h. As mentioned in Section 2.2.2, much of the data is now publically available. Regarding forecasts, only analysis data ("nowcasts" for hour zero) can however be accessed

Table 2.1. *Information on a selection of global, freely available wind forecast datasets. Horizontal and temporal resolutions are given for forecasts up to one week ahead.*

	GEFS	ECMWF	NCEP
Type	Reforecast	Forecast	Forecast
Horizontal resolution	0.5°	0.28°	0.7°
Time (from)	Dec 1984	Oct 2006	Mar 2007
Horizon (days)	0–15	0–14	0–15
Temporal resolution	3–6 h	6 h	6 h
Perturbed members	10	50	20
Runs per day	1	2	4
Height	80 m	10 m	10 m

free of charge. Furthermore, forecasts are not stored, so even if the policy would change, historical forecasts would not be available.

TIGGE, THORPEX Interactive Grand Global Ensemble, was a part of the international research collaboration THORPEX (The Observing System Research and Predictability Experiment). The aim with TIGGE was to provide operational ensemble forecast data to the research community [74]. Although the THORPEX project was finalised in year 2014, global forecasts from several weather services can still be accessed from e.g. ECMWF's web page[9].

NCEP's GEFS (Global Ensemble Forecasting System) [75] is, to my knowledge, the only long-term, global and freely available dataset of reforecasts. In analogue to reanalyses being produced from historical data using approximately consistent models and assimilation systems, reforecasts that are consistent with an operating forecast system can be produced.

In Table 2.1, information on GEFS second generation reforecast and the global forecasts from ECMWF and NCEP is given. The latter two were chosen out of the TIGGE members since these were most promising for our needs; the datasets should preferably span several decades, be consistent in time, have a high temporal and spatial resolution, have long enough forecast horizon, be issued several times each day, give wind forecasts at roughly the hub height of modern WTs and have several members (if probabilistic forecasts are desirable). As an overall judgement, GEFS was considered the most appropriate choice.

[9] `http://apps.ecmwf.int/datasets/data/tigge/levtype=sfc/type=cf/` Accessed: 2016-04-14.

3. Theory and Methods

In this chapter, a theoretical background for the papers included in the thesis is given. Most of the methods used in the papers are thoroughly described in textbooks on e.g. wind engineering, statistics and signal processing and references are given for further reading. Since Matlab has been the main tool for the modelling, references are given to both built-in functions and code available on Matlab's file exchange area. For readers familiar with modelling of wind power, parts of this chapter may seem trivial and can be omitted. Based on feedback on our work, I nevertheless believe that it can be useful to introduce the theoretical background and methods by starting from a more basic level.

3.1 Wind speed

Large-scale wind patterns are driven by pressure differences and the rotation of the earth (the Coriolis effect). On the meso- and microscale, phenomena such as sea breeze, terrain etc. also affects the wind flow. In more complex terrain, it can therefore be large differences in wind speeds on the scale of hundreds of meters. In the surface layer (up to about 100 m above ground [76]), the wind speed generally increases with height, i.e. there is a positive wind shear. The shear depends on surface roughness and orography. In and near forests, the wind speed is typically low near the ground, but increases sharply with height. The shear also varies in time depending on the atmospheric stability and other factors. The wind profile is sometimes described by the empirical power law:

$$u(z) = u_{ref}\left(\frac{z-d}{z_{ref}-d}\right)^{\alpha}, \tag{3.1}$$

where u is the wind speed, z is the height above ground, subscript *ref* denotes a reference height where a measurement is available, d is the (zero-plane) displacement height and α is the shear exponent. The displacement height describes the elevation of the zero level of the wind in and near cities, forest and other vegetation and is often around 0.6–1.0 of the canopy height. In practice, the wind speed never reaches zero, but the profile at heights relevant for wind power is better described if the displacement height is taken into account. Wind speeds at a particular site are often considered to follow a Weibull distribution,

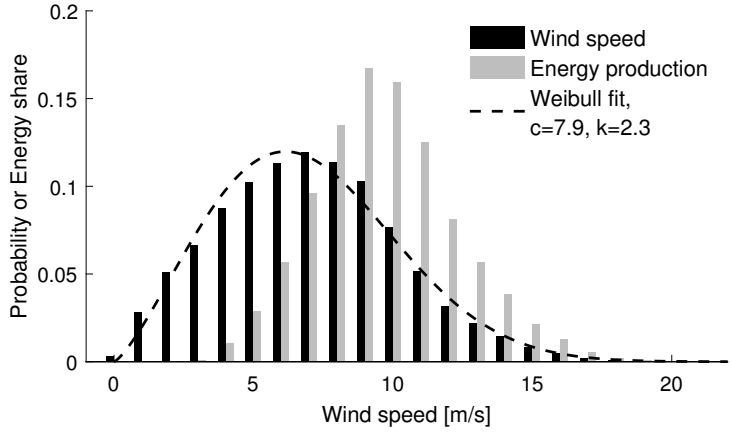

Figure 3.1. Histogram and Weibull fit of wind speeds for a measurement taken 100 m above ground (mean wind speed is 7.0 m/s). Energy production shares (see Section 3.2), calculated for Vestas V90 2MW, are shown in grey.

$$p(u;k,A) = \frac{k}{A}\left(\frac{u}{A}\right)^{k-1} \cdot e^{-(u/A)^k}, \tag{3.2}$$

where p is the probability density function of wind speed u, given the Weibull scale factor A and shape factor k. The Weibull parameters A and k are linked to the mean wind speed $\bar{u}$; given two of these parameters, the third can be calculated. If nothing more than mean wind speed is known, a Rayleigh distribution is commonly assumed, i.e. a Weibull distribution with shape factor equal to 2. One must though be careful assuming Weibull/Rayleigh distributed wind; some sites exhibit two peaks in the distribution and an assumed Rayleigh distribution can give severe errors in energy calculations. An example of the distribution of measured wind speeds and the corresponding Weibull fit is shown in Fig. 3.1.

The temporal variation in the wind speed is significant. Seasonal and diurnal patterns are present in many regions. In Sweden, the generation is in average 74% higher in November–March than in June–August, see Section 4.2.1. An example of wind speed variations for a measurement taken 96 m above ground is shown in Fig. 3.2.

In order to determine the wind conditions for a potential wind farm, on-site wind measurements are normally carried out during a one- or two-year period. Because of the variability of wind speed and direction, a long-term correction (LTC) of the measurement is necessary. Commonly, this is done with the "MCP" (Measure Correlate Predict) method [69]. A simple LTC method is described in Section 3.2.

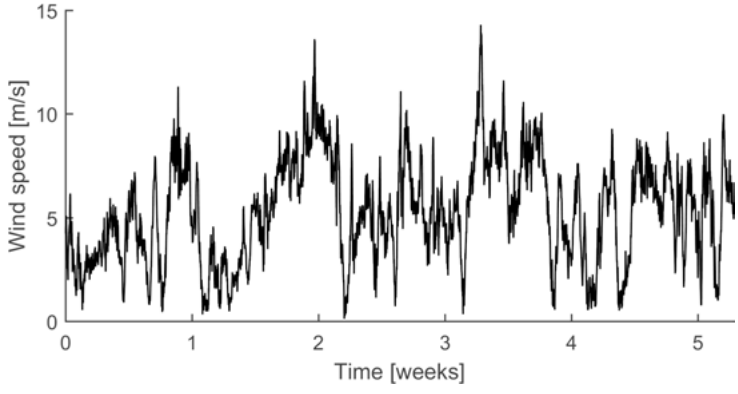

Figure 3.2. A few weeks wind speed measurements (10-minute resolution, 96 m above ground), showing the significant temporal variability.

3.2 Wind power

The available power P_w in the air mass flow is a function of the wind speed u, air density ρ and area A perpendicular to the wind direction:

$$P_w = \frac{1}{2}\rho A u^3. \tag{3.3}$$

It has been shown that the amount of wind energy convertible to mechanical energy has a theoretical upper limit of $16/27 \approx 59\%$ [77]. A real WT has a lower aerodynamic efficiency and by including mechanical and electrical losses the overall efficiency is often peaking at 45–50% for large-scale modern WTs. Because of component and grid connection costs, there is a trade-off in the rating of the generator. For a given turbine area, a generator with a higher rating will increase the electricity production but also the costs. The optimum ratio of rated power to rotor area (specific rating) depends on the site-specific wind conditions. The resulting power to wind speed relation is known as the power curve (PC) and can readily be obtained for any commercial WT, at least for standard air density (1.225 kg/m^3). According to the IEC standard [78], the PC is given as a function of 10-minute mean wind speeds. Some examples of normalized PCs are shown in Fig. 3.3. Three wind speeds of importance for the WT are the cut-in (where the WT begins to operate, often 3–4 m/s), the rated (the lowest wind speed with full output, often 11–14 m/s) and cut-out (where the turbine is shut down for protection, often 25 m/s). As can be seen in the example in Fig. 3.1, the largest share of the energy is often produced at relatively high, but less probable, wind speeds.

In reality, there are several factors influencing the power output of a WT. The PC accounts for aerodynamic and electrical losses in the WT. Additional reductions in the power occur due to aerodynamic degradation from dirt and

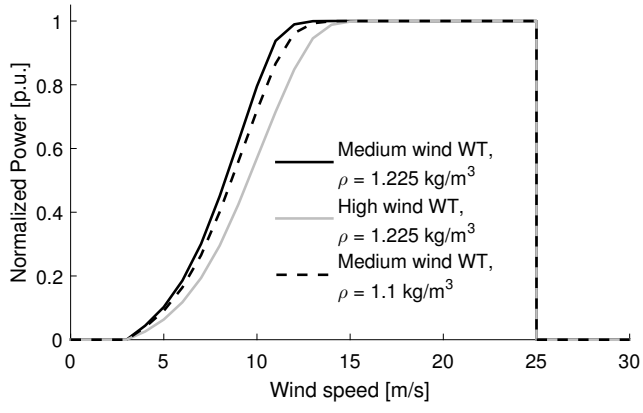

Figure 3.3. Three selected power curves for illustrating the impacts from specific rating and air density. "High wind" windturbine (WTs) have higher specific ratings and thus requires higher wind speeds to reach the rated power.

icing, hysteresis losses[1], grid and transformer losses and stops due to maintenance, failures or icing. Since the certification of the PC is performed under certain standard conditions, operation of the WT in a different wind regime might change the PC. As an example, a more turbulent flow gives higher power at low wind speeds but lower power near rated wind speed.

As described in Section 3.1, a LTC is routinely used to determine the long-term wind conditions at a site before the investment decision for a farm is taken. A basic LTC method called *wind index* [79] was used to determine the long-term corrected yearly averages of wind generation in Sweden, see Section 4.2.1. For each farm, a 35-year generation time series was computed from MERRA data which was first linearly scaled so that the average generation equalled that specified by the farm owner. The wind index (WI) for a farm i and year y was subsequently computed as

$$WI_{i,y} = \frac{\overline{P}_{i,y}}{\overline{P}_i}, \tag{3.4}$$

where $\overline{P}_i$ is the average power for the 35-year period and $\overline{P}_{i,y}$ the average power for year y. The national WI was subsequently calculated as the capacity-weighted average of the WIs of all farms in operation during the year in question. By dividing the measured annual generation by the corresponding wind index, the "normal year" generation is obtained.

[1] After a high wind stop the WT does not start until the wind speed falls below a certain value (which is lower than the cut-out wind speed).

3.2.1 Modelling

Most papers included in this thesis rely heavily on the idea of modelling aggregated wind power generation from reanalysis data. The basic model is presented in Paper I and extensions in Papers II and III.

Parameters were tuned to give good agreement between model output and measurements on the system scale. As the following example illustrates, this approach can have a strong impact on optimal parameter values. For a single farm, the output varies between zero and (very close to) rated capacity. If parameters were to be tuned to measurements for separate farms, the model would produce time series with a similar range. However, when such time series are combined, one would get nation-wide generation with almost the same range, which is clearly not correct; the yearly maximum of generation for the whole of Sweden is often around 80% of the installed capacity. One explanation to this phenomenon is that out of thousands of WTs in a country, some will always be out of operation. Even more important is the underestimation of wind speed diversity by coarse NWP models; we simply never have situations were all WTs simultaneously experience wind speeds above rated. Somewhat simplified, the methodology involves the steps in the numbered list below. The first six steps are described in detail in Paper I while step 7 is described in Paper II and step 8 in Paper III.

1. A database of WTs is created, containing e.g. coordinates, rated capacities and rotor diameters.
2. Time series of wind speed and direction for each WT are bilinearly interpolated[2] from MERRA reanalysis data.
3. The wind speed time series are scaled in order to get annual generation in agreement with that specified by the farm owner. As described in Section 2.2.1 this is, in my opinion, a better method than calculating the wind speed directly from reanalysis wind speeds at different heights.
4. Hourly generation time series are computed for each WT, taking into account farm specific PCs, losses of different kind, PC smoothing (see Section 3.2.4), wind direction etc.
5. The output is aggregated for the studied area and subsequently multiplied with a variable less than one ("external" losses, see Paper I).
6. Corrections are performed, accounting for seasonal and diurnal bias as well as the fact that the annual generation predicted by turbine owners was systematically too high in the 90's and early 00's.
7. Fluctuations in the high-frequency range, $f > (10h)^{-1}$, were underestimated by the model described in Paper I (points 1–6 above), see explanation below. This leads to an underestimation of step changes, which is not a good thing for power system studies. Thus a statistical model

[2] Vector summation is employed in order to get correct interpolation of wind direction (the mean of e.g. 359° and 1° should be 0° rather than 180°).

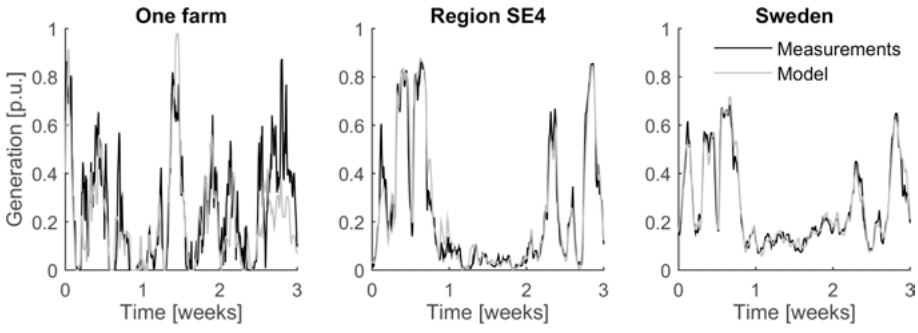

Figure 3.4. Time series of measurements and model outputs for a farm, a region and the whole of Sweden. It is clear that the two latter can adequately be modelled with relatively coarse reanalysis data, but for a single farm the results are not satisfactory.

was developed in order to increase these fluctuations, see Paper II and Section 4.1.2.

8. A statistical model, simulating intra-hourly variations, was developed in Paper III. A short methodology description is given in Section 4.1.3. This step is only executed if data with a higher resolution than hourly is necessary.

When NWP reanalyses are used as input for generating wind power time series, the high-frequency fluctuations are often not properly represented. For one WT, the magnitudes of these fluctuations are often underestimated. An explanation to this effect is the limited spatial and temporal resolution of the NWP models; local features causing variations in the wind speed are not captured. A model with a certain resolution is only able to correctly represent phenomena on significantly larger scales, see [33] for a more in-depth discussion. The "effective" resolution of NWP models is often in the order of seven times the nominal resolution. When the outputs from many farms in a power system are combined, the results can change; in [45] it was found that the variability in terms of step changes was severely overestimated by NWP models. This could be explained e.g. by too high correlations between the time series for different farms. In most studies, the high-frequency variability is nevertheless underestimated also for the aggregated generation, see Paper II.

A legitimate question is whether a relatively coarse reanalysis is appropriate for modelling wind power. As the results in Section 4.1 show, the answer to this question is yes. The following example however demonstrates that our approach is only suitable for modelling generation aggregated over relatively large areas, where errors for individual farms are cancelled out. Let us consider the output from one farm, one region in Sweden (SE4, see Fig. 1.1 on page 14) and the whole of Sweden. Fig. 3.4 clearly shows that the coarse NWP model cannot resolve the local wind variations seen by a farm but works fine for larger regions.

3.2.2 Scenarios

When modelling the output of future farms in a system, it is necessary to develop scenarios of deployment, both in terms of geographical locations and of technical specifications of the farms. Several scenarios of future Swedish farms were developed in Paper D. As discussed in Section 2.2.1, coarse reanalyses are neither suitable for determining likely sites for deployment of wind farms nor for estimating mean wind speeds. Even a more highly resolved NWP can be misleading for determining likely sites since many other factors than the mean wind speed have a strong influence. Distance to a suitable grid-connection point, quality of the infrastructure, the risk of icing, conflicting interests (natural values, the military, aviation, telecommunications etc.) and, not the least, the attitude of the local municipality, both neighbours and the local government, are often more important than the wind resource itself. Instead of trying to model all these factors, we started from a database of present and planned farms, see Fig. 4.15 on page 83. Since the cost for obtaining permits, measuring and analysing the wind etc. is generally several hundred thousand euros for each farm, it is reasonable to assume that such database gives a good estimate of the actual distribution of future WTs. The farms in the database were assigned probability weights depending on e.g. the status of the permitting process, mean wind speed according to the MIUU model (see Section 2.2.2) and whether the farm is located in a designated area for wind power or not. The technical characteristics of the farms, most importantly CFs, were estimated from MIUU wind speeds, historical trends and interviews with persons in the wind energy community.

In Papers VI and VII, scenarios of deployment in the Nordic synchronous power system and Europe were developed. Because of the larger scopes of these studies and since the authors do not have the same in-depth knowledge of wind power in these countries, simpler methodologies were used. Still though, the geographical locations were not based on coarse reanalysis wind speeds, but current and planned farms and/or mesoscale wind maps. Also, the mean wind speeds were not taken directly from reanalyses, but were rather computed from the anticipated energy yields and WT characteristics.

The appearance of the aggregated generation profile of a region or country naturally depends on both geographical distribution and technical specifications of the farms. In Paper D it was found that the dispersion of farms and average CF were the two most influential factors and consequently needed most attention for the development of scenarios. The CF for a farm is a function of e.g. mean wind speed, specific rating, wake losses and efficiency of the WTs. The mean wind speed, in its turn, is determined by the wind conditions at the site and the hub height. From a variability point of view, it is of relatively little importance how the average CF is achieved. As an example, farms with relatively low wind speeds and large rotor areas give more or less the same results as farms with higher wind speeds and smaller rotors.

3.2.3 Quantifying variability

The output from a WT varies on all time scales, from sub-seconds to decades. As mentioned in the introduction, the variability can be challenging for the power system since the fluctuations need to be balanced in one way or another. In order to quantify and visualise the variability, many different metrics and figures have been used in the literature. In Paper C, a review on variability metrics, these were classified as "Distribution", "Change" or "Temporal Characteristics". The same categorisation is used in the following exposition, where some of the metrics are presented and, when considered necessary, explained.

The distribution metrics naturally describe the distribution of the time series and are independent of observation order. The mean generation in its normalised form is often denoted capacity or load factor. A 3 MW turbine that produces in average 1.1 MW thus has a CF of $1.1/3 \approx 0.37$. The mean generation can also be expressed as equivalent full-load hours (or only full-load hours): the number of hours in a year that the turbine would have to run at its rated capacity in order to produce its annual energy. The abovementioned WT thus generates $0.37 \cdot 8760 \approx 3200$ full-load hours, which is typical for a good onshore project in Sweden as of today.

The minimum generation for a single WT or farm is zero, but also for a relatively large region (e.g. the Nordic synchronous area) the minimum is often very close to zero. I therefore believe that the minimum generation level is not a very informative metric; it is preferable to study low generation in terms of quantiles/percentiles[3], perhaps conditioned on the electric load. Maximum generation, in particular for larger regions, can however be useful to analyse. The firm capacity is defined as the output exceeded a certain percentage p of the time (e.g. 90%), and is thus the same as the $100 - p$ percentile.

A duration curve is a plot of all data sorted from highest to lowest and can often be an important tool for studying variability. Of particular relevance is the study of net load duration curves for different scenarios of wind deployment. An empirical cumulative distribution functions (ECDF) can be attained by a 90 degrees clockwise rotation of the duration curve. These plots thus hold the same information and it is only a matter of taste which to use. In most wind studies, duration curves are the preferred option. The sample standard deviation (SD) or variance can quantify the variability of the generation with one single number. Probability density function estimates in the form of histograms give more information.

Of equal or higher importance than the analysis of the distribution of generation is that of changes in generation, which can be measured as step changes or ramp rates. The former is defined as change of averaged power over a fixed number of steps, e.g. change in hourly energy over four hours. The definition

[3]See `http://se.mathworks.com/help/stats/quantiles-and-percentiles.html` for how these metrics are computed in Matlab.

of ramp rates, however, differs between authors. As put in [80] "*depending on the defining criteria of a ramp event, the starting and ending point and duration of a ramp can change substantially*". Some authors also define ramp rates in the same manner as step changes [81]. In our papers, one and four hour step changes in hourly generation are normally considered.

Most of the metrics used for quantifying the distribution of generation can be directly transferred to the analysis of step changes. The SD of step changes is possibly the most common of all variability metrics. Since the distribution of step changes can be skewed and normally have fat tails, the higher moments (skewness and kurtosis) are sometime given [37]. It can be mentioned that the changes in power depend on the initial generation level. Sörensen *et. al.* [82], among others, therefore plotted percentiles of ramp rates over initial power.

In Paper C, metrics in the "Temporal Characteristics" category, admittedly a bit confusing, included:

- Metrics of temporal characteristics of the data (excluding the metrics of change discussed earlier)
- Frequency domain metrics

The latter type of metrics is treated in Section 3.4. The autocorrelation function (ACF) is the linear correlation of a time series with itself, shifted in time. If, for instance, the ACF is high for lag one, it implies that the generation is not likely to change much from one sample point to the next. In some cases, e.g. for wind power forecast errors, a high ACF is not desired since it implies that the energy imbalance (actual - forecasted generation) can build up to significant levels. The sample ACF at lag k of a stationary time series $x_1, ..., x_n$ can be computed as

$$ACF_k = \frac{\sum_{t=k+1}^{n}(x_t - \bar{x})(x_{t-k} - \bar{x}))}{\sum_{t=1}^{n}(x_t - \bar{x})^2}. \tag{3.5}$$

The volatility of wind power is not constant; sometimes the output is relatively steady and sometimes it fluctuates heavily. A useful way of quantifying the varying volatility is the moving variance, which is the variance of a signal in a moving window[4]. The moving SD is defined analogously. When the moving SD is applied on a filtered time series, only the fluctuations of certain frequencies are considered. This metric, referred to in [34] as the "Variability Index", was used in Papers II and III.

3.2.4 Smoothing

When the outputs from several WTs are combined, be it in a farm or in a whole synchronous power system, the resulting time series is smoother than that for

[4]A fast implementation in Matlab (*movingvar*) has been written by Aslak Grinsted, see `http://www.mathworks.com/matlabcentral/fileexchange/8252-moving-variance` Accessed: 2016-08-04.

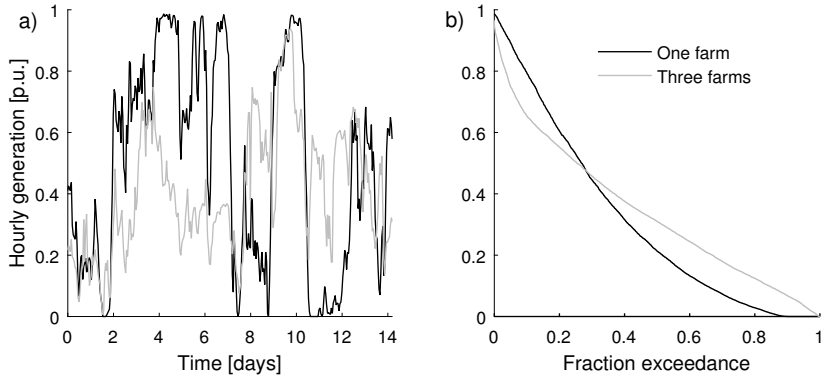

Figure 3.5. Smoothing of output when combining several farms. **(a)** Two weeks example of time series. The profile for three farms combined is smoother; ramp rates are lower and the generation is more seldom very high or very low. **(b)** Duration curves based on 19 months of measurements.

a single WT [6]. An example of outputs from one farm and three farms in aggregation is given in Fig. 3.5. The three farms combined have lower ramp rates and the generation is more seldom very high or very low.

If the correlations between generation time series are known, the reduction in variability (in terms of variance) can be quantified. Let y_t be a weighted sum of n discrete time series $x_{i,t}$

$$y_t = \sum_{i=1}^{n} w_i x_{i,t}. \tag{3.6}$$

The variance of y can then be computed from the weights w_i, the standard deviation σ_i of each $x_{i,t}$ and the pair-wise, linear correlations ρ_{ij} between the time series:

$$Var(y_t) = \sum_{i=1}^{n} \sum_{j=1}^{n} w_i w_j \sigma_i \sigma_j \rho_{ij}. \tag{3.7}$$

As an example, let us consider the aggregation of two time series $x_{1,t}$ and $x_{2,t}$, each with $\sigma = 1$ and $w = 0.5$. With ρ_{12} of -1, 0, and 1, the variance of the combined outputs become 0, 0.5 and 1 respectively. Eq. (3.7) is, due to its low computational cost, very powerful for optimising the w's for lowest possible variance of y_t.

The smoothing effect makes it beneficial to connect many farms over a wide area. The variability is reduced and hence integration of wind power becomes less of a challenge. As was demonstrated in Eq. (3.7), a low correlation between outputs gives better smoothing. A commonly [36] used model for describing the correlation between farms is given by

$$\rho(d) = \alpha \cdot e^{-d/D}, \tag{3.8}$$

where d is the distance between the farms and D is a parameter determining how fast ρ decreases with distance. The α parameter is sometimes fixed to unity and sometimes allowed to take other values as well. In Paper VI, the correlation coefficients between wind generation in different countries in Europe were quantified and it was shown that Eq. (3.8) is appropriate also for country-wise generation.

There is also smoothing within a farm or a smaller region. A "multi-turbine" PC therefore gives a better representation of the output than the original PC from the manufacturer. A methodology for calculating this is proposed in [83]. In our database (see Section 2.1), detailed coordinates are not available for all WTs. Furthermore, no turbulence intensity data was available and the temporal resolution of the wind speed time series was one hour instead of the ten minutes which is the standard for PCs. Thus a simplified version of the methodology given in [83] was employed, using only one smoothing parameter which was tuned to generation data for the whole of Sweden. Let us assume there are several WTs of the same type distributed within one grid cell in the meteorological model. At a certain hour, the mean wind speed is 6 m/s which, according to the manufacturers PC, translates to a generation of 0.19 p.u. However, some WTs will experience lower or higher mean wind speeds than 6 m/s and there will also be differences within the hour for each WT. We therefore assume, as in [83], that the 10-minute wind speeds for all WTs are normally distributed around 6 m/s. The smoothing parameter was defined as the standard deviation of this distribution. In our example, a smoothing parameter of 2 m/s results in an average generation of 0.24 p.u for the considered hour. An example of an original PC and its multi-turbine counterpart is shown in Fig. 3.6. For low and above cut-out wind speeds, the output is higher for the multi-turbine PC. For wind speeds between a few m/s lower than rated and cut-out, the output is lower. In our implementation, the same smoothing parameter was applied for all WTs and all wind conditions. Although this is a quite rough simplification, a smoothing parameter of around 1.0 m/s improved the model of aggregated Swedish wind generation considerably as compared to using original PCs, see Fig. 11 in Paper I.

3.2.5 Forecasts

For secure operation of power systems and profitable operation of generators, load and intermittent generation need to be predicted in advance. Forecasts for different time horizons are of interest, e.g. short-term forecasting for intra-day operation, day-ahead forecasts for energy trading and one-week forecasts for hydropower planning. For more information, in particular on Sweden and the Nordic system, see [84].

A real example of generation, day-ahead forecasts and forecast errors (FEs) for a wind farm is given in Fig. 3.7. The forecast generally follows the gen-

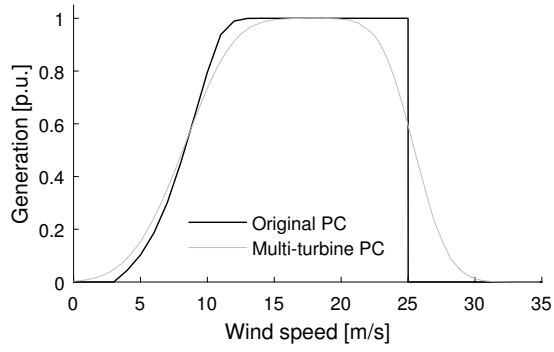

Figure 3.6. Examples of original and multi-turbine power curves (PCs). The latter was calculated with a smoothing parameter of 2 m/s (see Section 3.2.4).

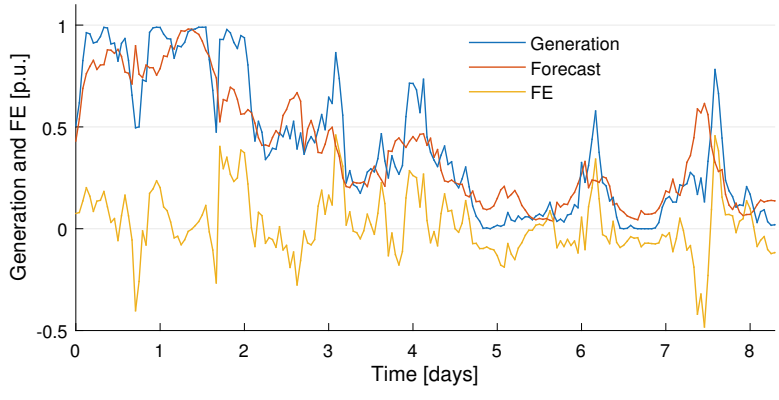

Figure 3.7. Example of actual generation, day-ahead point forecasts and forecast errors (FE) for a farm in Sweden. Note the phase error at day seven, leading to first a large negative error and then a large positive error.

eration quite well, but for some hours the mismatch is large. Two main types of FEs exist: level and phase errors [85]. The latter (misjudgements of when the generation ramps up or down) are short-lived but can be very large. The FEs are autocorrelated; if the error is positive one hour it is likely that it is also positive the following hours.

Traditionally, research focus has been on improving the performance of point forecasts [86] but recently a lot of attention has been drawn to probabilistic forecasts, providing the user with prediction intervals. Commonly, probabilistic forecasts are obtained from ensemble NWP models where several members are produced by perturbing the initial conditions. Different NWP models can also be combined, both to improve point forecasts and to get information on prediction intervals [87].

An example of an ensemble forecast from the GEFS dataset is shown in Fig. 3.8. For this particular forecast, there are relatively small differences

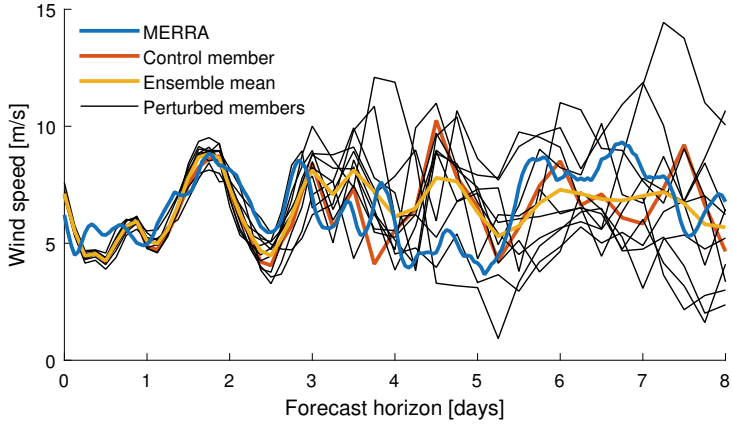

Figure 3.8. Example of an ensemble forecast from GEFS. The ten perturbed members are obtained by perturbing the initial conditions of the control forecast.[6]

between the ensemble members for the first three days, but for longer horizons the spread is larger. The control member, shown in blue, is the forecast for unperturbed initial conditions.

For low penetrations of wind power, the uncertainty in net load is dominated by load uncertainty. Wind power FEs however become increasingly important for higher wind penetration levels and thus need to be accounted for in wind integration studies. According to IEA wind [18], several years of "synthetic" forecasts are preferable for simulation of unit commitment and economic dispatch including reserve requirements. As for generation time series, synthetic forecasts can stem from (linear scaling of) historical data, NWP models, statistical models or hybrids of the latter two. A review of wind integration studies published between 2007 and 2016 (see Paper IV), revealed that NWP based forecasts were rarely used but purely statistical approaches were common.

Actual FEs have several characteristics, some of which make them challenging to simulate with purely statistical methods: the errors increase with terrain complexity [88], depend on the wind speed [89] and are often skewed and heavy-tailed [90]. Furthermore, there exist level and phase errors and consecutive forecasts are not independent[7].

Despite the challenges, many attempts have been made to simulate realistic wind power forecasts using statistical methods. After some initial experimenting, we decided to use a different strategy: to base the synthetic forecasts on GEFS reforecasts (see Section 2.2.3) and some statistical fine-tuning. In or-

[6]Note the butterfly effect; small changes in the initial conditions of a chaotic system eventually leads to large differences. This is why weather is so difficult to predict for longer time horizons (for all but Laplace's Demon).

[7]If, for example, the five day-ahead (D+5) forecast issued a certain day is too conservative (forecast < generation), then it is likely that the D+4 forecast issued the following day is also too conservative.

der to retain the benefits of NWP based forecasts (from a modelling perspective), e.g. the occurrence of phase errors and dependencies between consecutive forecasts, we strived to keep the statistical post-processing to a minimum. In short, the method involves the following steps, also illustrated in Fig. 3.9:

1. The "true" generation is calculated from reanalysis data with the methodology described in Section 3.2.1.
2. "Raw" wind speed forecasts for day 0 through 7 are interpolated from GEFS reforecasts.
3. All or most of the bias in the forecasted mean wind speed is removed.
4. Some multivariate ARMA noise, described in Section 3.3.4, is added to the forecast in order to decrease the correlation between FEs for different farms and obtain an appropriate spectrum.
5. A fixed part of the errors for each hour is removed.
6. Generation forecasts are computed from wind speed forecasts.
7. A machine learning model (Gradient boosting, see Section 3.6.3) is optionally used in order to improve the generation forecasts.

For trimming the model, measurements/forecasts from the whole of Belgium and individual farms in Sweden were used. FE correlations were also compared to results in the literature. Although no probabilistic forecasts were generated, this could be done in a straight-forward manner by using the perturbed members of GEFS.

3.3 Statistical methods

This section comprises five subsections. First, different methods for optimisation and the partitioning of data into training, validation and test sets are reviewed. Secondly, the transformation between different distributions is described briefly. The three last subsections are devoted to splines, ARMA models and empirical orthogonal functions respectively.

3.3.1 Optimisation

In the papers included in this thesis, several different optimisation methods were employed for finding suitable parameter sets. In general, the main objective was not to infer "true" parameter values or find causal relationships, but rather to build models suitable for prediction. As discussed in [91], there are often multiple models (parameter sets) having similar performance in this sense.

In most models there are two types of parameters that need to be optimised: hyperparameters and model parameters. The former determines the structure of the model and are set beforehand. As an example, in classification and

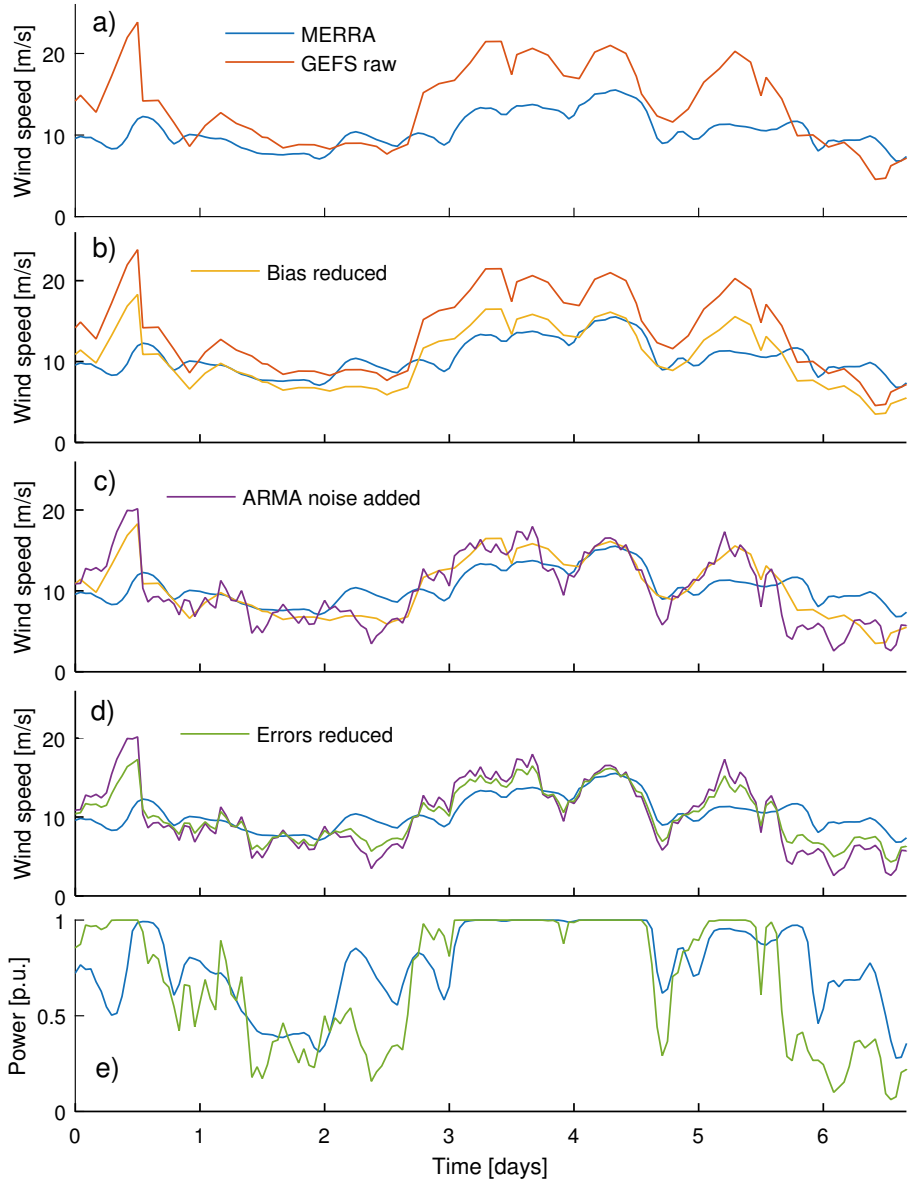

Figure 3.9. Illustration of the method for generating synthetic forecasts. The steps referred to in the following are described in the numbered list on page 43. **(a)** Step 1–2, **(b)** Step 3, **(c)** Step 4, **(d)** Step 5, **(e)** Step 6.

regression trees, the maximum number of splits is a hyperparameter set by the user and the actual splits are determined by the data when the model is trained. Different algorithms are often employed for optimising the two types of parameters. E.g. a grid of hyperparameters can be searched and for each set of hyperparameters, the splits that minimise a loss function are chosen. In most of my papers, the function that should be minimised is denoted objective function. An objective function can be of the type reward/utility function (if it should be maximised) or loss function (if it should be minimised).

When training and evaluating the performance of a model, it is crucial to not use the same data for these two purposes. If many different sets of hyperparameters are considered, it is useful to partition the data into three independent sets:

1. **Training data** for determining optimal model parameters given certain hyperparameters.
2. **Validation data** for determining optimal hyperparameters. For each set of hyperparameters, a model is trained using the training data and the performance is assessed with the validation data.
3. **Test data** to evaluate the performance of the chosen model.

We can illustrate the concepts of optimisation with a simple example. Let

$$Y(X) = 40 + 0.8X - 0.01X^2 + 0.00002X^3 + \varepsilon,$$
$$X \sim \mathscr{U}(0,500) \quad \varepsilon \sim \mathscr{N}(0,100^2), \quad (3.9)$$

where $\mathscr{U}(0,500)$ is the continous uniform distribution (equal probability in the range 0–500 and zero probability outside this range), be an unknown function we try to fit with polynomials

$$P_n(X) = a_0 + a_1X + a_2X^2 + ... + a_nX^n. \quad (3.10)$$

In this case n is a hyperparameter that is determined beforehand. All values of n between one and nine were considered. Three datasets were generated with $N = 100$ sample points each: training, validation and test sets. The model parameters (coefficients) $a_0, ..., a_n$ were obtained by minimising the squared residuals of the training data. Fig. 3.10 show some important characteristics of the optimisation:

- The data is somewhat scattered. In reality this can be due to e.g. measurement uncertainty and neglected predictors.
- Model performance is generally best for the training data.
- For the training data, a more complex model (higher n is this case) gives a better fit. For $n \geq N - 1$, the polynomial would pass through all data points.

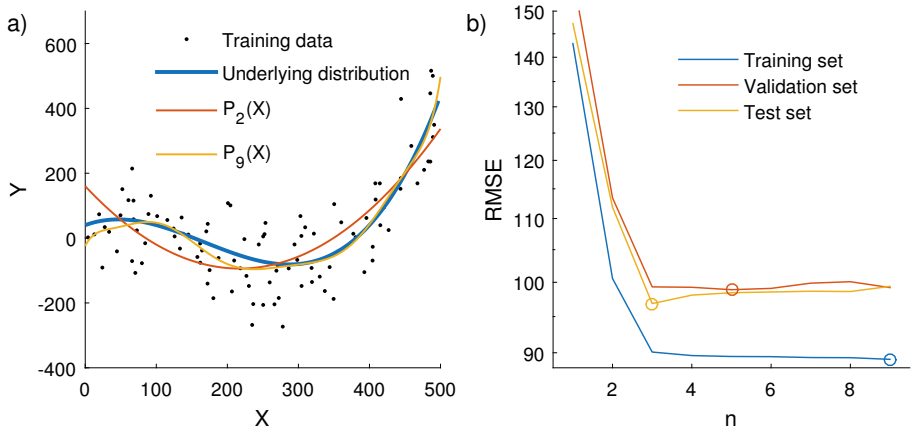

Figure 3.10. **(a)** Training dataset, underlying distribution $Y(X) = 40 + 0.8X - 0.01X^2 + 0.00002X^3$ and two examples of polynomial fits. **(b)** RMS error for the different sets depending on polynomial order n. The best polynomial orders for each set are indicated with circles.

- Due to overfitting, a too complex model will give poorer results for the validation and test data. The $n \geq N-1$ polynomial mentioned in the previous point would most likely give terribly bad results for the validation and test sets.
- If the same dataset is used for determining optimal values of the hyperparameters and assessing model performance, overoptimistic results may be achieved. In our example, $n = 5$ is optimal for the validation set which results in a higher RMS error for the test set as compared to $n = 3$.

The simplest method for optimisation is probably the search over a regular grid. For each parameter, allowed values are specified and subsequently the objective function for all possible combinations of parameters is evaluated. This method was employed e.g. in Paper II, where a coarse grid of a few hyperparameters was evaluated. For high-dimensional problems and/or when a fine resolution near local optima is desired, this method quickly leads to impractically long computational times.

In Paper I, a random restart hill-climb algorithm was used for optimising the model of hourly wind generation in Sweden. This algorithm starts with a random set of parameter values and the performance of the model is evaluated. Subsequently, a randomly selected parameter is increased or decreased. If this leads to a better model, the change is accepted and a new randomly selected parameter is altered. In the beginning of the search, relatively large changes of the parameter values were used but subsequently smaller steps in the parameter space were taken. When no further improvement is possible, the parameter set is saved and a new hill-climb with a random start point is undertaken. In Paper I, the model with the lowest objective function for the training set was

chosen. In hindsight, this led to a slightly overfitted model; somewhat better results could have been attained if a validation set had also been used.

Matlab's built-in *fminsearch*, utilised in several of our papers, employs the Nelder-Mead simplex algorithm [92, 93] to find the minimum of a function. A simplex in n dimensions has $n+1$ vertices, e.g. a triangle in two dimensions. The function is evaluated at the vertices. New, better simplexes are constructed by replacing one of the vertices, e.g. by reflection of the worst point through the centroid of the remaining points. This unconstrained, derivative-free algorithm often requires relatively few function evaluations, although it can sometimes fail to find the true minimum [94]. A transformation can be used to convert a bound constrained problem into an unconstrained problem. This has been implemented in the *fminsearchbnd* script[8].

If the variance of a combined time series is to be minimised, e.g. optimising the allocation of wind power capacity in several countries in order to get a smooth output, *fminsearchbnd* with Eq. (3.7) as loss function can be used for very fast computations (a few seconds). Naturally, zero is set as the lower boundary for each weight, but additional constraints may also be needed. In Papers VI and VII, the variance was e.g. minimised while keeping the total energy production constant.

In two of the papers, phase angles of the frequency domain representation of signals were optimised. The algorithm for finding appropriate phases is described in Section 3.4.2.

3.3.2 Transformation between distributions

Transformations between distributions were used in Paper III for obtaining normally distributed, stationary time series. The transformations were made either from a t location-scale distribution (see Paper III) fitted to data to the normal distribution. Since it is not possible to express the cumulative distribution functions of these distributions in terms of elementary functions, the transformations could not be written as closed-form expressions. Instead, the Matlab *cdf* and *icdf* commands were used. For fitting a distribution to data at hand, the Matlab *fitdist* function was used. An example of the fitting and transformation procedure is shown in Fig. 3.11. First a t location-scale distribution is fitted to the 1000 sample points. The dashed line illustrates how a sample point is then transformed to the normal distribution ($-0.005 \rightarrow -1.18$).

3.3.3 Splines

It is often desirable to smoothen data in one way or another. An option to low-pass filtering (described in Section 3.5) is to use splines; piecewise polynomi-

[8]Script written by John D'Errico, `http://www.mathworks.com/matlabcentral/fileexchange/8277-fminsearchbnd--fminsearchcon` Accessed: 2016-08-04.

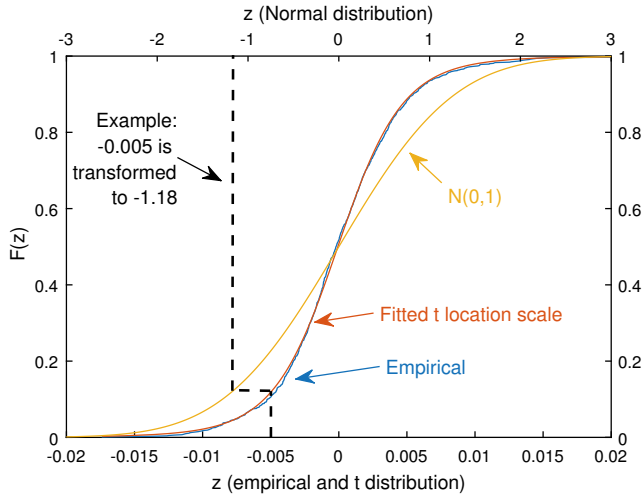

Figure 3.11. Illustration of the transformation from a fitted t location-scale to a normal distribution. The vertical axis gives the cumulative distribution functions $F(z)$. Note that two different horizontal axes are used and that the t location-scale distribution gives a good fit to the empirical data.

als with a continuous derivative. In Paper III, wind power fluctuations with temporal resolution of 5–15 minutes were simulated and added to smoothened hourly generation. Because of the limitation in article length, the smoothing splines used, both Matlab's built-in cubic splines and our own fourth order splines, were described relatively briefly. More details are thus given in the following paragraphs. Let us first define the following variables:

f_s	Sampling frequency (samples per hour)
P_t	Mean generation with frequency f_s
P_h	Mean hourly generation
C_t	Cubic spline fit of P_h, evaluated at resolution f_s
x_t	High-frequency component ($P_t - C_t$ for the training data)

Let also additional subscripts *sim0*, *sim* and *obj* denote preliminary simulated time series, simulated time series and the objective power system respectively. For generating $C_{t,obj}$ from $P_{h,obj}$, Matlab *fit* command with *fitType* set to *cubicinterp* was used. The algorithm fits cubic polynomials between each pair of data points. Cubic interpolation is a special case of smoothing splines where the fitted function passes through all data points. If desired, smoothing splines can be made smoother, at the expense of not passing through all data points, by adding a penalty for the integral of the squared second order derivative of the splines.

The preliminary simulated generation was achieved as $P_{t,sim0} = C_{t,obj} + x_{t,sim0}$. Due to the stochastic nature of $x_{t,sim0}$, an adjustment was necessary

in order to make the mean of $P_{t,sim0}(t \in h)$ the same as $P_{h,obj}$. Let ΔP_h denote the difference in hourly means for hour h. Simply adding ΔP_h to $P_{t,sim0}$ would sometimes result in large step changes at new hours. For the f_s samples in hour h, a fourth order spline adjustment $S_h(r)$ was therefore used:

$$x_{t,sim}(t \in h) = x_{t,sim0}(t \in h) + S_h(\boldsymbol{n}), \tag{3.11}$$

where

$$S_h(\boldsymbol{n}) = a_h \cdot \boldsymbol{n}^4 + b_h \cdot \boldsymbol{n}^3 + c_h \cdot \boldsymbol{n}^2 + d_h \cdot \boldsymbol{n} + e_h \tag{3.12}$$

and **n** is the vector (0 1 ... f_s-1). The coefficients $a_h, ..., e_h$ in Eq. (3.12) were chosen to satisfy

$$\begin{cases} S_h(0) = \frac{\Delta P_{h-1} + \Delta P_h}{2} \\ S_h(f_s) = \frac{\Delta P_h + \Delta P_{h+1}}{2} \\ \overline{S}_h(\boldsymbol{n}) = \Delta P_h \\ S'_h(0) = S'_{h-1}(f_s) \\ \text{Minimise} \int_0^{f_s} (S''_h)^2 dr \end{cases} \tag{3.13}$$

The coefficients $a_h, ..., e_h$ are uniquely defined by the constraints in Eq. (3.13). The algebraic expressions for the coefficients can be derived in a straightforward fashion but are quite long, so they are not given here. $P_{t,sim}$ was subsequently calculated as

$$P_{t,sim} = C_{t,obj} + x_{t,sim}. \tag{3.14}$$

In an earlier version of the adjustment function, cubic splines were used and the last condition in Eq. (3.13) was not included. This resulted in oscillations with a one-hour period, clearly observable as a peak in the power spectral density (see Section 3.4.1) of $P_{t,sim}$. Examples of the difference in hourly mean generation $\Delta P_h(t)$ as well as the spline adjustment $S_h(t)$ and the resulting $P_{t,sim}$ are shown in Fig. 3.12.

3.3.4 Multivariate ARMA models

In Papers II-III, IFFT method was utilized to generate noise with a desired frequency domain spectrum, see Section 3.4.2. This is a powerful method in that it allows synthesis of time series with an arbitrary spectrum. In Paper IV, we desired noise with certain correlations between different sites. An Autoregressive Moving Average (ARMA) model was thus deemed more appropriate since the correlations can then be controlled.

An ARMA process with AR order p and MA order q is notated ARMA(p,q) and satisfies the equation

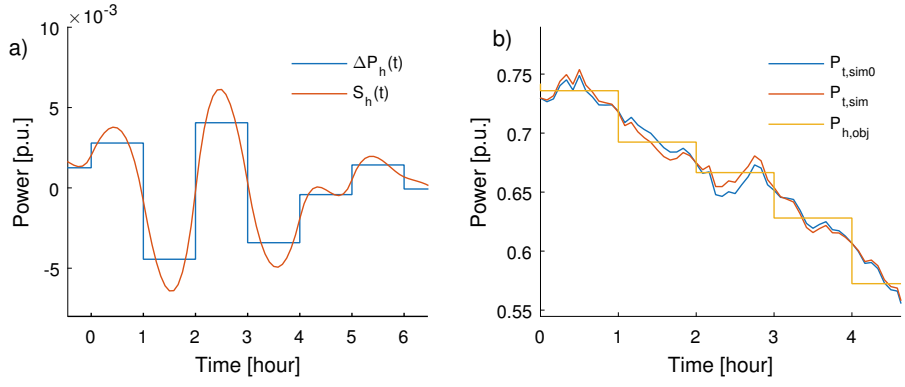

Figure 3.12. **(a)** Spline adjustment for correcting differences in mean hourly generation (ΔP_h). If ΔP_h would be added directly, abrupt step changes at new hours would occur. **(b)** Resulting time series $P_{t,sim}$ after adjustment. Note that the same hours are not shown in the two sub-figures.

$$Y_t = \phi_1 Y_{t-1} + \phi_2 Y_{t-2} + ... + \phi_p Y_{t-p} + \varepsilon_t + \theta_1 \varepsilon_{t-1} + \theta_2 \varepsilon_{t-2} + ... + \theta_q \varepsilon_{t-q}, \tag{3.15}$$

where ϕ_i are AR coefficients, θ_j are MA coefficients and ε is Gaussian noise with mean zero and variance σ^2. In order to generate correlated Y_t time series, multivariate ε noise were generated with the Matlab *mvnrnd* function. The entry for row n and column m in the covariance matrix, necessary as input to *mvnrnd*, was computed as

$$c_{nm} = \sigma^2 e^{-d_{nm}/D}, \tag{3.16}$$

where d_{nm} is the distance between wind farms n and m and D is a decay factor. This technique resembles that used in e.g. [95]. As an example, let $D = 50$ km, $\sigma = 2$ m/s and $d_{1,2} = 17$ km. This gives $c_{1,2} = 2.85$ and the ε time series for farms 1 and 2 will thus have a correlation coefficient close to $2.85/2^2 = 0.71$. This correlation is preserved in the ARMA filtering. The generation of noise from the ARMA model was implemented with a Matlab filter using a rational transfer function determined by the ARMA parameters, which in turn were fitted to give an appropriate spectrum.

3.3.5 Empirical orthogonal functions

Empirical Orthogonal Functions (EOFs) are described relatively informally in this section, see e.g. [96, 97] for a more rigorous treatment. Meteorological variables, e.g. temperature, wind speed and pressure, are normally available as time series for each point in a grid. The detrended (mean removed) data for a variable can be ordered in a matrix $\mathbf{X}$:

$$\mathbf{X} = \begin{bmatrix} x_1(1) & x_2(1) & \cdots & x_m(1) \\ x_1(2) & x_2(2) & \cdots & x_m(2) \\ \vdots & \vdots & \vdots & \vdots \\ x_1(n) & x_2(n) & \cdots & x_m(n) \end{bmatrix}, \tag{3.17}$$

with n time steps and m grid points. For a highly resolved model and/or large geographical area, m quickly becomes large and it may be inappropriate or impossible to use all m time series as regressors in a model. Fortunately, an EOF decomposition allow us to reduce the dimensionality of the problem significantly with very little loss of information. Informally, one finds the dominating spatial patterns of the variable (the EOFs) and the corresponding EOF time series. Often, most of the variance of the variable van be explained by just a handful EOFs. Mathematically, a decomposition of a real-valued $\mathbf{X}$ into three matrices is performed with singular value decomposition:

$$\mathbf{X} = \mathbf{U} \cdot \mathbf{D} \cdot \mathbf{V}^{\mathbf{T}} \tag{3.18}$$

where $\mathbf{V}$ contains the EOFs and $\mathbf{U} \cdot \mathbf{D}$ are the EOF time series. In Papers II and IV, EOFs were computed in Matlab using the *pcatool*[9]. Note that different terminology is sometimes used in the literature; EOFs can e.g. be denoted principal components and the EOF time series denoted principal component amplitudes. Apart from reducing the dimensionality, EOFs can also reveal important characteristics of a system in terms of modes of variability.

We can illustrate the use of EOFs by looking at wind speeds from the MERRA reanalysis dataset for Sweden (latitude $54-70^\circ$ and longitude $10.0-25.3^\circ$). The horizontal resolution is $1/2^\circ$ in latitude and $2/3^\circ$ in longitude, so $m = 33 \times 24 = 792$. Fig. 3.13 shows the first three EOFs[10]. In Fig. 3.14, the actual wind speeds for an hour is compared to the sum of mean wind speeds and the first nine EOFs for the same hour. As can be seen, the patterns are very similar. For the whole dataset, the first nine EOFs explain 90% of the variance and the first 20 EOFs explain 99%. The 792 original time series can thus, for regression purposes, be replaced by e.g. nine or twenty time series. Note that the number of EOFs needed can differ significantly; for some variables used in Paper II, the first EOF alone explained over 90% of the variance and for some, 14–15 EOFs were needed. Also note that in an EOF analysis, one has to determine the geographic region to include, see e.g. [35].

[9]Script written by Guillaume Maze, available at `http://www.mathworks.com/matlabcentral/fileexchange/17915-pcatool` Accessed: 2016-08-04.

[10]Note that each EOF corresponds to a column in $\mathbf{V}$. For presentation, each column vector was first reshaped to the native grid of MERRA.

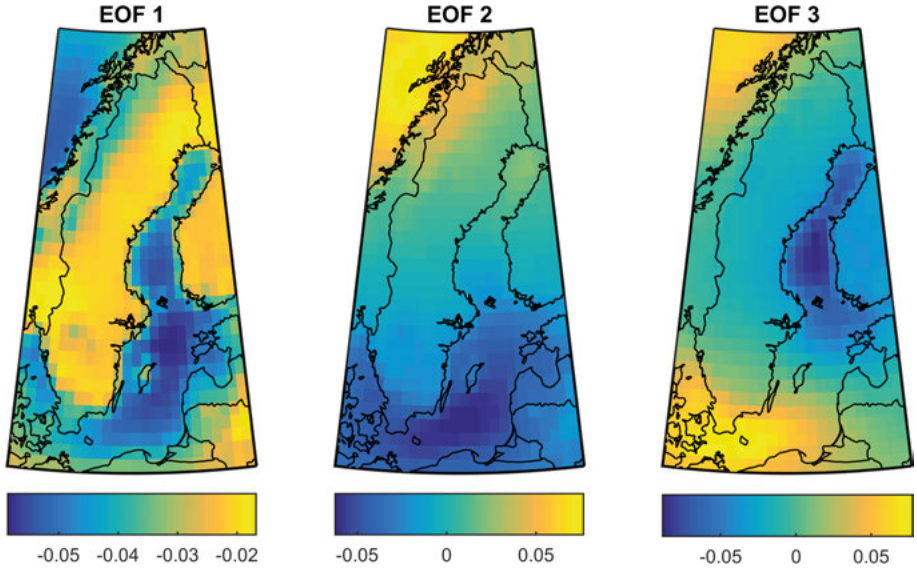

Figure 3.13. The first three EOFs of wind speed. EOF number one is strictly negative and strongly negatively correlated to the mean wind speed (-0.94). A positive contribution of this EOF thus gives lower and more even wind speeds. The second and third EOFs imply certain patterns predominantly in the offshore winds.

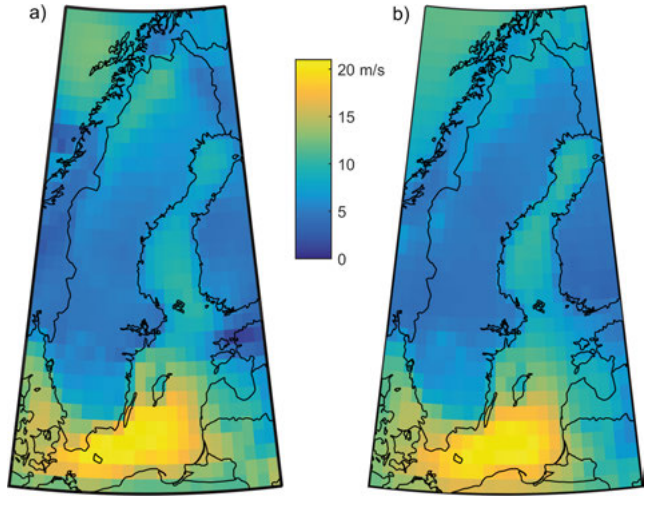

Figure 3.14. Illustration of the use of EOFs. **(a)** MERRA wind speeds for a certain hour and **(b)** sum of mean wind speeds and the first nine EOFs for the same hour. These EOFs explain 90% of the variance. For regression purposes, the 792 original time series, one for each grid cell in **(a)**, can be replaced with only e.g. nine.

3.4 Frequency domain methods

A discrete signal x_t in the time domain can be decomposed into its frequency domain components X_k by the fast Fourier transform (FFT). Each element in X_k can expressed in rectangular notation ($a_k + b_k i$) or in polar notation (magnitude and phase). The inverse procedure (IFFT) synthesises a time domain signal from the frequency domain components.

3.4.1 Power spectral density

In many signal processing applications, an important aspect of a FFT analysis is to isolate and study the dominating frequencies in a noisy signal. For wind power, the frequencies corresponding to one year and one day (and their harmonics) are often prominent. Wind power generation is however, to a large degree, stochastic and all frequencies are of interest. A useful method to analyse the strength of fluctuations with different frequencies is power spectral density (PSD) estimates, which are closely related to the squared magnitudes of the frequency-domain representation of the signal. The simplest PSD estimate comes from the periodogram and is computed as

$$PSD_k = \frac{1}{f_s \cdot N} |X_k|^2 \tag{3.19}$$

The periodogram is often very noisy. In order to reduce the variance of the PSD estimate, several different options are available. The Welch's PSD estimate [98] split the signal in several overlapping segments. Each segment is windowed; samples near the ends are reduced in amplitude (in the Matlab *pwelch* function, the Hamming window is the default). This is done in order to reduce spectral leakage, see e.g. [99]. The final PSD estimate is calculated as the average of periodograms of the windowed segments. Two disadvantages with segmenting the data are that the frequency resolution becomes poorer and that information on very low frequencies is lost. If very long time series are available, it is possible to achieve both high resolution and low noise. An option is to compute the Welch PSD for high frequencies using multiple segments and the PSD for low frequencies with one or a few segments.

Log-log plots of the Welch PSD estimates for Swedish wind power generation as well as output from one farm are shown in Fig. 3.15a. Some characteristics, typical for wind power PSDs, are worth mentioning:

- The PSD is generally lower for higher frequencies (shorter periods). The contribution from relatively high frequencies to the variance of the signal can however be significant, see Fig. 3.15b.
- The PSD for the whole of Sweden is considerably lower than that for one farm when the generation is given in p.u., indicating that the combined output is smoother than that for one farm.

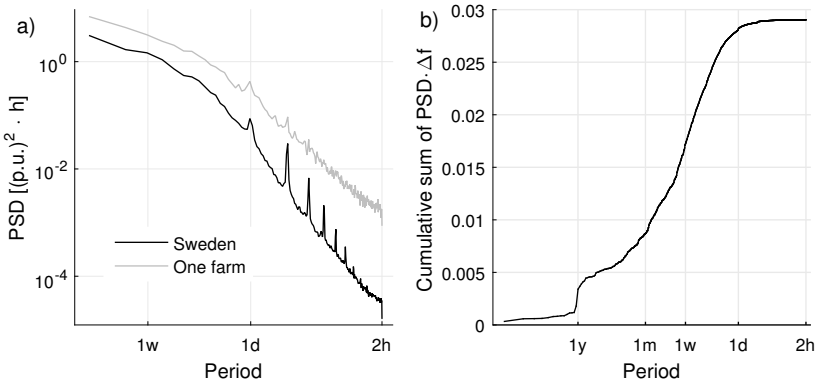

Figure 3.15. Power spectral density (PSD) estimates of hourly wind generation. **(a)** Welch estimates with 257 sample segments for the whole of Sweden and one farm, **(b)** cumulative sum of $PSD \cdot \Delta f$ computed with periodogram for the whole of Sweden. The contributions sum up to the variance of the time series (0.029 p.u.), see Eq. (3.20).

- The PSD slope is steeper for the whole power system, see also [100, 101]. This is a result of higher frequency components of different farms being less correlated than lower ditto and thus more effectively dampened by aggregation, see Paper VI.
- There is a smaller peak corresponding to diurnal wind variations. Because of the non-sinusoidal diurnal patterns, harmonics with frequency $(24h/2)^{-1}, (24h/3)^{-1}, \ldots$ are also present.

In Fig. 3.15b the cumulative sum of $PSD \cdot \Delta f$ for the whole of Sweden is shown (Δf is the frequency increment). An advantage with this type of plot is that noise in the PSD becomes less visible, thus a periodogram of the whole time series can be used and information on lower frequencies becomes available. If the mean value of a signal x_t is subtracted before the PSD is computed, as is the case for Fig. 3.15, we furthermore have

$$\sum_k PSD_k(x_t) \cdot \Delta f = Var(x_t). \tag{3.20}$$

Fig. 3.15b can thus be interpreted as the contribution to the overall variance of the generation from different frequency bands; a steep curve in a certain region implies that this is important in terms of variability. Very little of the variance can be explained by variations with periods shorter than 24h (around 3%). Almost half (47%) of the variance can be attributed to fluctuations with periods between two days and two weeks while around 9% comes from the yearly cycle.

PSD analyses were used in Papers II, III, IV and VII. Generally, the frequencies were expressed in terms of periods since this is easier to interpret, e.g. $f = (24h)^{-1}$ instead of $f \approx 0.417h^{-1}$. The Hz unit was seldom used since

this becomes unpractical for the time scales considered. As mentioned above, it can be advantageous to remove the mean from the signal before computing the PSD. If this is not done, the relation in Eq. (3.20) does not hold and the PSD for the lowest frequencies can be distorted. The mean was not always removed in the papers if it was not important for the analysis.

3.4.2 Synthesis of time series

In Papers II and III, time series of wind generation were constructed by combining low-frequency components from meteorological models and stochastic high-frequency noise. This was done either because no high-frequency data were available (Paper III) or because the high-frequency fluctuations from the models were too smooth. Other authors have used statistical methods to generate purely synthetic time series of generation and forecasts [31, 44, 101–104], but as discussed in Section 1.3, we believe that it is beneficial to use physical models as far as possible.

For the following exposition, let P_t denote the wind generation in discrete time and x_t the high-pass filtered P_t. Let also an additional subscript *sim* indicates simulated data and no additional subscript indicates measurements. The volatility of x_t vary considerably in time. In Papers II and III, x_t was therefore transformed to an approximately stationary time series y_t, see Sections 3.3.2 and 3.5.

When synthesising $y_{t,sim}$, there are often several criteria that need to be fulfilled e.g. appropriate distribution, autocorrelation, spectrum and step change distribution. Furthermore, the varying volatility of x_t should preferably be captured. The FFT-IFFT approach provides a very flexible method that, together with e.g. stationarity transformations and predictions of volatility, can meet these demands. The analysis and synthesis (FFT and IFFT) themselves are not complicated, the challenging part is to find appropriate magnitudes and phases of Y_{sim}, the frequency domain representation of $y_{t,sim}$. The magnitudes of Y_{sim} were interpolated from the magnitudes of Y_t. In an early version of the methodology (Paper D), uniform random phase angles were used to synthesize $y_{t,sim}$. This gives a correct PSD, but $y_{t,sim}$ becomes normally distributed, which was not always desired. Furthermore, the moving SD of $y_{t,sim}$ fluctuated quite a bit which led to a deterioration of the models' ability to predict the varying volatility of x_t. Inspired by [101], we therefore used an algorithm to search for appropriate phases.

Let τ_t denote the moving SD of $y_{t,sim}$ and σ_τ the SD of τ_t. The loss function L to be minimized was chosen as

$$L = \sigma_\tau + c \cdot |\kappa_{sim} - \kappa_{desired}|, \tag{3.21}$$

where κ is the kurtosis and c is a constant. The latter term in Eq. (3.21) was included in order to impose appropriate heaviness of the tails on $y_{t,sim}$. In

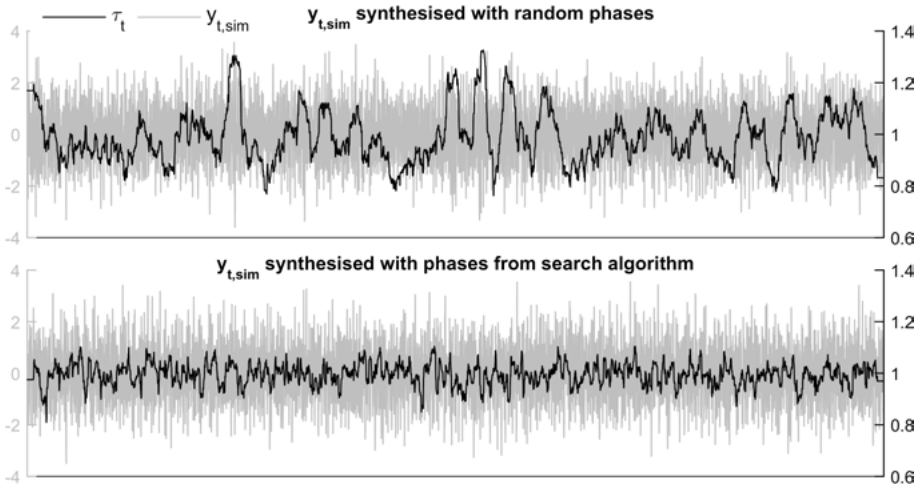

Figure 3.16. The volatility of τ_t (the moving standard deviation of $y_{t,sim}$) is reduced by using a phase angle search algorithm before synthesising $y_{t,sim}$. Based on data from Paper III.

order to minimize L, it was found that an "elitist" algorithm with small but many generations was robust and efficient:

1. Start with ten uniform random phase angle sets (ranging from $-\pi$ to π).
2. Synthesize time series (a generation with ten individuals) from the magnitudes and phases using IFFT and calculate L.
3. Choose the angles corresponding to the individual with lowest L. Construct ten new angle sets by adding random noise from a $\mathscr{U}(-a,a)$ distribution.

Step 2 and 3 were repeated 2000 times with the a parameter gradually reduced from $\pi/10$ to $\pi/300$. The reduction in the volatility of τ_t for one case in Paper III is illustrated in Fig. 3.16; σ_τ was reduced from 0.104 to 0.041. A further reduction would be possible, but 2000 iterations was considered a reasonable trade-off between low σ_τ, desired kurtosis and computational time.

The computational time for simulating a 300,000 sample point time series was around ten minutes, which should be compared to only ten seconds with random phases. The improved performance of the model must therefore be weighted against the increased computational time.

Two comments can be made on the methodology outlined above. Firstly, it is often the case that the length of the simulated time series is different from the training data. As an example, one might have a few years of training data and want to simulate high-frequency fluctuations for a 30+ year long time series stemming from reanalysis data. In this case, the magnitudes of Y must be interpolated and scaled by $\sqrt{length(y_{t,sim})/length(y_t)}$ in order to get correct magnitudes for Y_{sim}. Secondly, if $y_{t,sim}$ should be real-valued, Y_{sim} must be

made symmetrical: even symmetry for the magnitudes and odd symmetry for the phases.

3.5 Filters

Filters of different kinds were used in almost all papers. For a very accessible introduction to filters and other signal processing techniques, see [99]. It should be recognised that the author of this thesis is not an expert on filter design; more knowledgeable persons at our division are gratefully acknowledged for their help on choosing appropriate filters. The motives for filtering data were e.g.

- Smoothing time series, e.g. with moving average filters.
- Separate the high- and low-frequency components of a signal, e.g. when intra-hourly wind power fluctuations were to be simulated.
- Splitting a signal into several frequency bands in order to deepen the analysis of variability and correlations.

Operations like ARMA simulations and calculations of step changes of wind power generation can also be interpreted as filtering of data. The kernel (impulse response) defines the filter in the time domain; the output from a filter is the input signal convoluted with the kernel. An example of a moving average (MA) filter is shown in Fig. 3.17. The kernel has the amplitude 1/10 for sample 5–14 and zero for all others. This means that e.g. sample 47 of the output signal is achieved as $\frac{1}{10}\sum_{i=5}^{14} x_{47-i}$, where x_n is sample n of the input signal. The length of the output from a convolution is the sum of the lengths of the input and of the kernel minus one. In order to get an output that is the same length as the input and without phase-shift, the following adjustment can be done. First, the kernel is made symmetrical, e.g. for a 5 point MA filter the kernel is defined as 1/5 for sample -2 through 2. Sample n in the output, y_n, is thus the average of samples $n-2$ to $n+2$ in the input. Secondly, the endpoints of the input signal need some special treatment. In the Matlab MA implementation (the *smooth* function), this is done by letting $y_1 = x_1$, $y_2 = (x_1 + x_2)/2$ etc. Zero-phase implementation and endpoints handling of other filters are described below.

Filters can be classified as infinite impulse response (IIR) or finite impulse response (FIR). For an IIR, the kernel has an infinite length, i.e. all samples in the output are affected by all samples in the input. An example of IIR is the low-pass sinc filter, which kernel is given by

$$h_n = 2 f_C \, sinc(2 f_C n), \tag{3.22}$$

where f_C is the cutoff frequency and $sinc(\cdot)$ is the normalised sinc function defined by

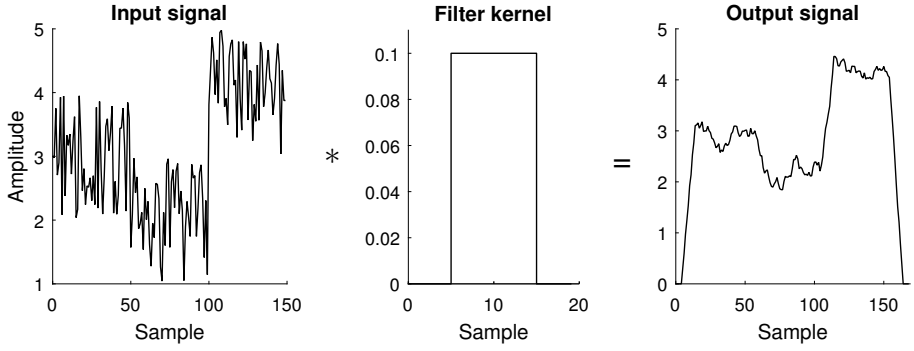

Figure 3.17. Moving average filtering of a noisy signal. The input signal is convoluted with the filter kernel, resulting in a smoother output signal.

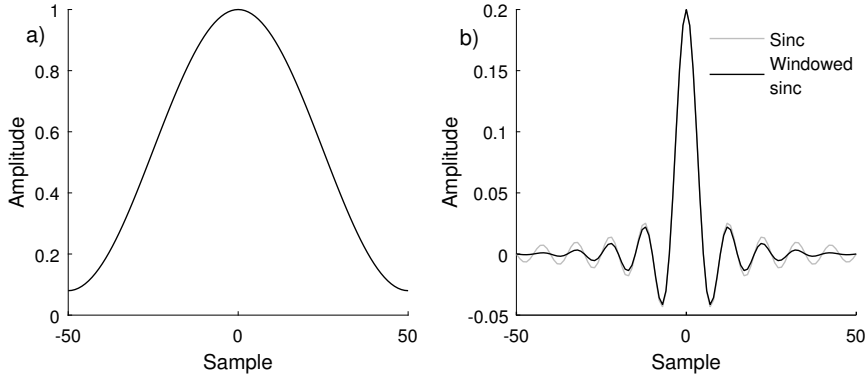

Figure 3.18. **(a)** Hamming window, **(b)** Kernels for a low-pass sinc and a 101 point windowed sinc filters with $f_C = 0.1$. The latter was windowed with the Hamming window shown in **(a)**.

$$sinc(x) = \begin{cases} 1 & x = 0 \\ \frac{sin(\pi x)}{\pi x} & \text{otherwise.} \end{cases} \tag{3.23}$$

Ideally, all frequencies lower than f_C are passed and all higher frequencies are blocked. When the sinc filter is windowed (i.e. the kernel is multiplied by a window function) it becomes a FIR filter called windowed sinc. An example of such filter kernel is shown in Fig. 3.18.

In Paper VII, high-order Butterworth filters were used. For numerical stability, these were implemented as second order sections (serially cascaded biquadratic sections). For both windowed sinc and Butterworth filters, two issues possibly need to be dealt with. Firstly, applying a filter will give a phase shift in the output. If the filtering is not performed in real-time, this issue can be resolved by applying the filter in both forward and reverse directions, e.g. using the Matlab *filtfilt* function. Secondly, the endpoints of the input time

series requires special attention, especially for low-pass filtering. In Paper VII this was handled by padding the time series with 10,000 samples of the mean of its first and last 1000 samples respectively. The resulting 20,000 extra samples of the output signal were subsequently removed. In Paper VI, no padding was performed but the first and last year of the filtered time series were simply removed.

Depending on the purpose of the filtering, different characteristics are more or less important. In signal processing, it can e.g. be desired to filter a square wave with noise. The time domain of the output is then in focus, e.g. an effective reduction of noise while keeping sharp edges of the square wave. For the applications in my papers, it was often more important to have sharp divisions in the frequency domain. The windowed sinc and high-order Butterworth filters both provide good separation of the frequency bands. Furthermore, when separating a signal x_t into frequency components $y_{1,t}, y_{2,t}, ...$, each with its distinct frequency content and together covering the whole frequency range, the following two relations hold to a very close approximation:

$$x_t = \sum_i y_{i,t} \tag{3.24}$$

$$Var(x_t) = \sum_i Var(y_{i,t}) \tag{3.25}$$

The latter equation is a direct result from Eq. (3.7) given that the different components are uncorrelated.

3.6 Machine learning

With the increased computational power, many different machine learning (or statistical learning) methods have become accessible. These often perform better and can handle more predictor variables than parametric procedures such as linear regression. The ability to handle many predictors makes machine learning suitable to use in combination with EOFs of meteorological variables, see Section 3.3.5. A common feature for machine learning techniques is that the functional forms linking the predictors to the response are derived inductively from the data and the techniques can thus be called non-parametric [105]. In Papers II and IV, random forests and/or gradient boosting were employed, both based on classification and regression trees (CART). Other methods include e.g. neural networks and support vector machines. An accessible introduction to machine learning with many references for further reading can be found in [105].

3.6.1 Classification and regression trees

CART [106] are the building blocks for many of the more advanced machine learning techniques, including random forests and gradient boosting presented in the following two subsections. As the name suggests, CART can be used either for classification of categorical response variables or for regression of quantitative response variables, and is particularly useful for high-dimensional problems. The focus of this and the forthcoming subsections will be on regression trees, since these were used in Papers II and IV.

CART works by recursively partitioning the data, i.e. earlier stages are not revisited after results from later stages are known. Assuming a quadratic loss function, the impurity i in node τ is defined by the within-node sum of squares:

$$i(\tau) = \sum (y_j - \bar{y}(\tau))^2, \tag{3.26}$$

where $\bar{y}(\tau)$ is the mean of all sample points in node τ. At each stage, the split that maximises the impurity reduction is chosen. [105] The partitioning process is illustrated in Figs. 3.19-3.20 with the example function

$$Y := sin(4 \cdot X) + Z^2 + \varepsilon, \quad X, Z \sim \mathscr{U}(0,1) \quad \varepsilon \sim \mathscr{U}(0,4), \tag{3.27}$$

i.e. Y is a non-linear function with two predictors X and Z and a stochastic component, but no interaction effects. As can be seen in Fig. 3.19, the values of Y, indicated by different colours, are generally lower when X is high. The best first split, reducing the impurity mostly, is for $X = 0.80$ (solid line). The second split (dashed line) was made at $Z = 0.66$ for $X < 0.80$. The partitioning can continue until there are very little data, sometimes only one sample point, in each terminal node. In order to avoid overfitting, it is often recommended to limit the number of partitions. This can be done by using a penalty function for tree complexity, limit the number of splits or to set a lower limit for the number of samples in the terminal nodes. Fig. 3.20 shows a grown tree with the maximum number of splits set to five. One can easily follow a path from the root node to a terminal node; the sample mean of Y for $X < 0.80$, $Z < 0.66$ and $X >= 0.19$ is for instance 2.92.

Regression trees can be used for prediction in a straightforward manner: a sample is dropped at the root node, and the predictors determine in which terminal node the sample ends up. The forecasted value of Y is the mean of the training data for that node. Although the CART methodology might seem simple, it is surprisingly good at handling multiple predictors and complex dependencies between these. Nevertheless, even better results can be accomplished by combining several trees, which will be the topic of the following two subsections.

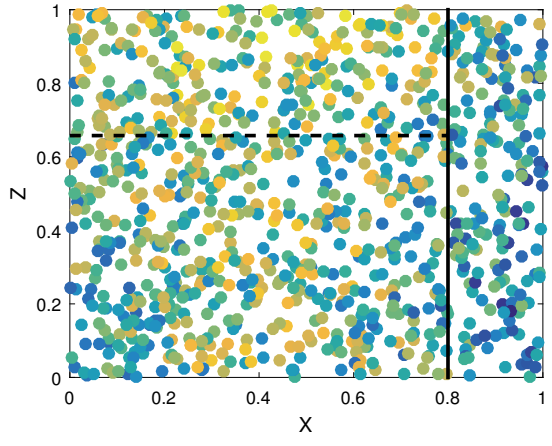

Figure 3.19. Partitioning of data with classification and regression trees (CART). Different colours indicate the value of $Y = f(X,Z)$; blue (low), yellow (high). The fist split is marked with a solid line and the second split with a dashed line.

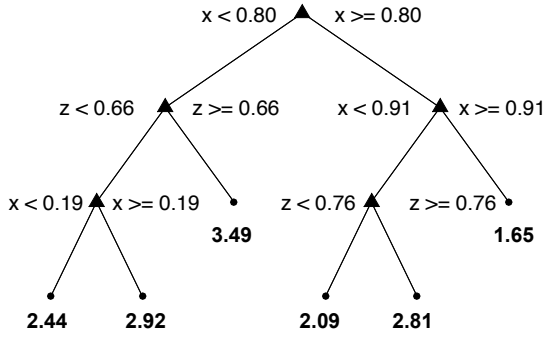

Figure 3.20. Regression tree with five splits and thus six terminal nodes (see Fig. 3.19 for the dataset that has been partitioned). The sample mean of Y for a certain region in the $X-Z$ space can be found by following a path from the root node to a terminal node.

3.6.2 Random forests

The idea behind random forests [107] is to grow an ensemble, often hundreds or thousands, of trees and combine the results in order to get more accurate and robust predictions. For each tree, a bootstrap sample[11] is taken from the data. For each partition, only a random sample of the predictors are considered. In regression mode, one third of the predictors are normally chosen for each partition. Generally, deep trees (many splits) are allowed to be grown. The data samples that were not sampled are called "out-of-bag" (OOB) and amount to around 37% of all data. The OOB data is subsequently dropped down the tree and mean values are computed for each terminal node. For predictions, the average of the predictions for each individual tree is used. The advantage of using OOB data is that the tree becomes robust to overfitting. One reason that the random forest algorithm is so effective is that only a subset of the predictors are used in each partition. This implies that very local features of the data will be captured in some of the trees. As a contrast, using only one CART with all regressors will only let the most dominant predictors influence the splits. [105]

In order to determine the importance of each predictor, Breiman [107] suggests to permute the predictors one by one and measure the increase in error, but other approaches are also possible [105]. It can often be useful to study the impact of each predictor. This can be done by partial dependence plots constructed in the following manner [105] for predictor X:

1. Grow a forest.
2. For all unique values of X, make a new database where other variables are left unchanged but X takes that value.
3. Plot the average predictions for each value of X.

Predictor importance and partial dependence plots for our earlier example in Section 3.6.1 is shown in Fig. 3.21. Matlab *TreeBagger* with standard settings was used to grow a forest with 100 trees (which is a relatively small forest). The Matlab *OOBPermutedVarDeltaError* property was employed to infer predictor importance. In the wind power field, random forests have been used by e.g. [34, 35, 108] and in Paper II.

3.6.3 Boosting

Boosting is also an ensemble method that can be based on CART, but there are some important differences to random forests. The latter, as we saw in Section 3.6.2, grow deep trees that are as independent as possible by using random samples of the data and the predictors. Boosting algorithms, on the other hand, improve the model by building upon the predecessor model. In

[11] A random sample with replacement of the same length as the data. A bootstrap sample of [1, 2, 3, 4] can for instance be [3, 1, 3, 4].

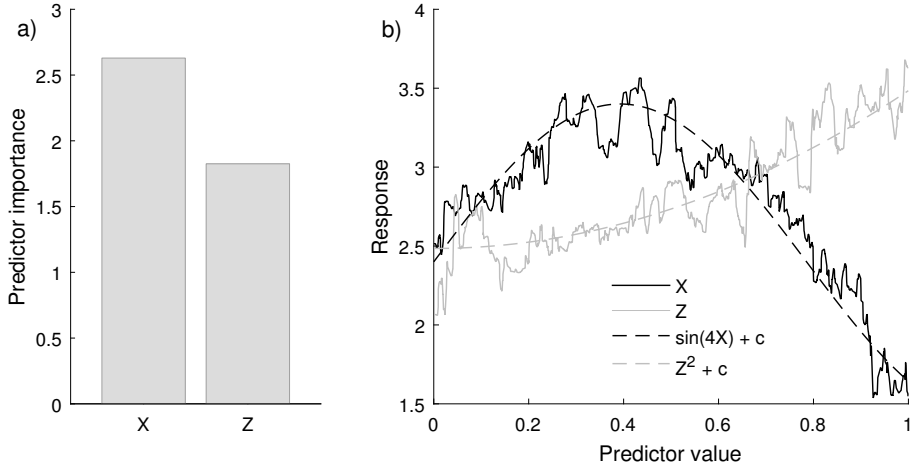

Figure 3.21. Predictor importance and partial dependence plots for a random forest trained on the dataset presented in Section 3.6.1 **(a)** Predictor importance in terms of increase in RMS error when variables are permuted **(b)** Partial dependence plots. In this example, the partial dependence plots are in good agreement with the underlying first-order contributions to Y (dashed lines).

contrast to random forests, there is generally no chance element involved in boosting algorithms and only a few splits are allowed for each tree.

Gradient boosting [109] (GB) stage-wise improves the model by finding and adding the tree that minimises the residuals from the previous tree:

$$F_m(\boldsymbol{x}) = F_{m-1}(\boldsymbol{x}) + \nu\rho_m h(\boldsymbol{x};\boldsymbol{a}_m), \tag{3.28}$$

where $F_m(\boldsymbol{x})$ is the model at iteration m and $h(\boldsymbol{x};\boldsymbol{a}_m)$, which is a CART with parameters $\boldsymbol{a}_m$, corresponds to the steepest gradient descent in the loss function. The optimal length of the descent is ρ_m, but by multiplying ρ_m with a *learning rate* parameter $\nu \in [0,1]$, the learning process can be forced to be slower, thereby increasing flexibility at the expense of an increased computational time.

When GB is used in regression mode minimising least square residuals, Friedman [109] calls it *LSBoost*. This algorithm is implemented in the Matlab *fitensemble* function. With random forests, there is a built-in protection for overfitting since OOB data are used for determining the fitted values for each terminal node. Boosting procedures generally does not have this built-in protection, so it may be useful to set aside some of the data as a validation set and use this to determine the optimum hyperparameters, e.g. number of iterations, learning rate and maximum number of splits for each individual CART. In the wind power field, boosting trees have been used by e.g. [35, 110] and in Papers II and IV.

4. Results and discussion

This section is structured in the following manner. First, in Section 4.1, evaluation results from modelling of aggregated wind generation and forecasts (Papers I–IV) are given, i.e. the model outputs are compared to measurements. Section 4.2 is devoted to historical and future wind generation in Sweden and contains results from Paper D and unpublished results based on datasets, models and methods described in this thesis. Section 4.3, finally, is based on Papers V-VII and contains results on variability and correlations of wind power and other IRE sources in the Nordic synchronous system and Europe.

4.1 Modelling wind power

Improving the methods for modelling wind power for power system studies is the main objective of this thesis. This section therefore contains the most important contributions. Firstly, results are given for the "basic" model of hourly, aggregated wind power generation (Paper I). Subsequently, results from Papers II and III are presented; improvements of the high-frequency part of the basic model and simulations of intra-hourly fluctuations respectively. Finally, the model for generating synthetic wind power forecasts is evaluated (Paper IV). The focus here will be on comparing model outputs to actual measurements and not so much on using the models for exploring the characteristics of future wind power. The latter is however the main motivation for developing the models, see Chapter 6 for some thoughts on future studies.

4.1.1 Basic model

The basic model, described in Paper I and briefly in Section 3.2.1, is based on MERRA reanalysis data and information on WTs. Separate models were trained for the whole of Sweden (SE) and bidding zones SE2, SE3 and SE4. Three years of data were used for training the models (2007, 2009 and 2011). The results given below are for the three evaluation years 2008, 2010 and 2012. No model was developed for SE1 since the data on operating farms in this zone was flawed: the measured generation was sometimes significantly higher than the installed capacity. During the studied time period, the installed capacity in SE increased from around 600 to 3500 MW (see Fig. 2 in Paper I). Especially in the first years, very few farms were operating in SE2 and the

Table 4.1. *Performance of the "basic" wind power model when comparing to measurements from the Swedish transmission system operator. Results are given for the evaluation years 2008, 2010 and 2012 for Sweden (SE) and three regions therein (SE2–4). All errors are given as percentages of the installed capacity.*

	SE	SE2	SE3	SE4
Mean Absolute Error	2.9%	6.5%	3.7%	4.2%
RMS Error	3.8%	9.1%	5.0%	5.9%
Mean Error	-0.1%	-0.7%	-0.5%	0.4%
Correlation	0.98	0.89	0.97	0.97

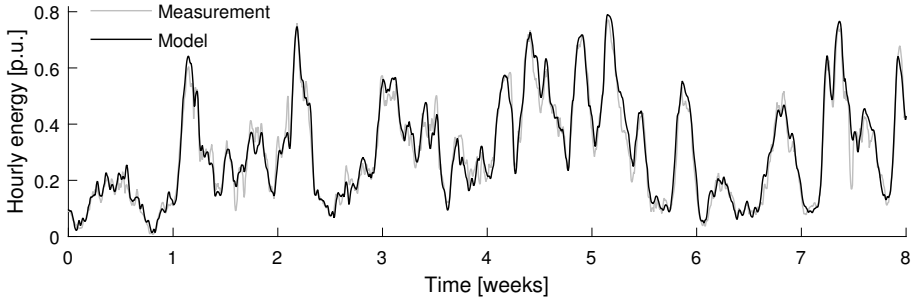

Figure 4.1. Evaluation of the modelling of hourly wind generation in Sweden (SE). Model output and validation data for eight weeks. The RMS error for the period is 0.039 p.u., i.e. almost the same as the average for all three years of validation.

installed capacity was below 100 MW. It is thus not surprising that the errors are higher for this region.

Some key performance metrics and results are given in Table 4.1 and Figs. 4.1–4.2. Overall, the results are satisfactory for SE, SE3 and SE4, but less so for SE2. For SE, the RMS error was 3.8% and the correlation to measurements 0.98, which is good in comparison to earlier results in the literature. As shown in Fig. 4.2, the distributions of hourly generation are also in good agreement. Before the application of a season/diurnal correction term, there was a systematic overprediction of wintertime generation by the model. A plausible explanation is icing losses.

Several different hyperparameters were tuned to give the best performance (low objective function). When the air density is higher, the available power in the wind increases. However, including air density in the computations of hourly energy did not improve the results. A possible explanation is the relationship between higher air density and icing losses. It was also found that constant wind shear coefficients for each farm gave as good results as using hourly values computed from the MERRA wind speeds at different heights. The inclusion of PC smoothing, seasonal/diurnal corrections and different losses in different wind sectors, in contrast, improved the model considerably.

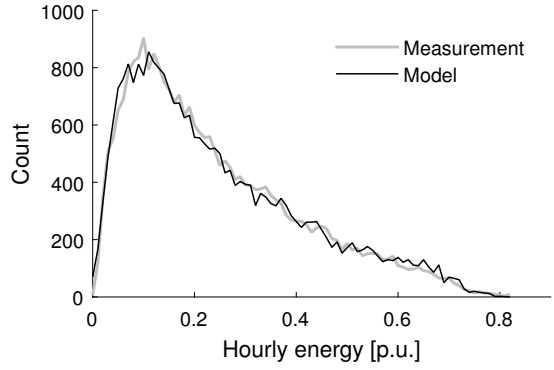

Figure 4.2. Histogram of hourly wind generation in Sweden (SE) calculated for bins of width 0.01 p.u.

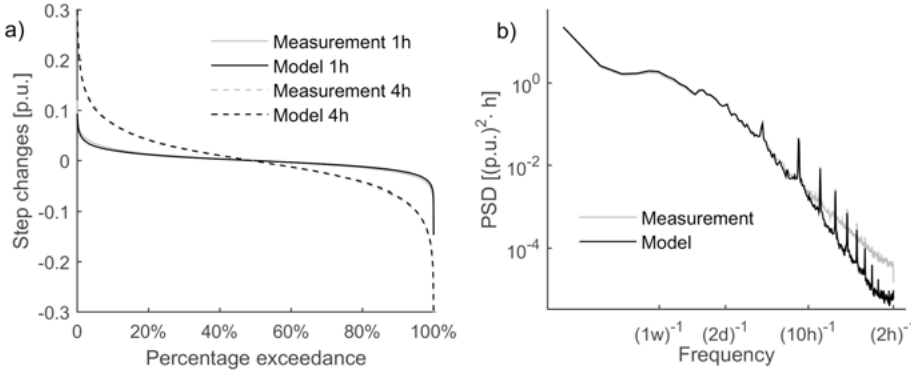

Figure 4.3. The basic model underestimated high-frequency fluctuations of wind generation. **(a)** Step change duration curves. The distribution of four hour step changes are adequately captured by the model, but one hour step changes are underestimated (see also Fig. 4.6). **(b)** Power spectral density (PSD) estimates. The PSDs starts to deviate for frequencies above $(10h)^{-1}$.

The output from the basic model was overly smooth which led to an underestimation of step changes. As an example, the SD of one hour step changes was understimated by 14% for SE and even more for the separate bidding zones. It is possible to detect the difference in smoothness in the time series plot in Fig. 4.1 and in one hour step change magnitudes in Fig. 4.3a. The difference is however much more obvious from PSD estimates, see Fig. 4.3b. The distribution of step changes can be crucial for determining the challenges of integrating more wind power into a power system. The lack of high-frequency variability was the reason for developing the improved model described in the next subsection.

4.1.2 Improved model

In this section, the model for simulating high-frequency variations of hourly, aggregated wind power generation is described briefly and results from a comparison with measurements are given. The elements in the modelling, e.g. filtering, FFT-IFFT transformations and machine learning algorithms, are described in more details in Chapter 3. As in Paper II, the basic model described in the previous subsection is denoted "the MERRA model", measurements are denoted "SvK" and subscript "sim" denotes simulated time series. The methodology involves the following steps:

1. Separate hourly time series from SvK and the MERRA model (M) into their low- and high-frequency components, e.g. SvK_{HF} and M_{LF}. A windowed sinc filter with a 5000 samples long kernel and cutoff frequency of $(10\text{h})^{-1}$ was used.
2. Transform SvK_{HF} to an approximately stationary time series y_{SvK}. This was done by dividing y_{SvK} by its variability index (VI), defined as the 24-hour moving SD of a signal.
3. Find the magnitudes of the frequency domain representation of y_{SvK} using FFT (Fast Fourier Transform).
4. Generate y_{sim} of the same length as the MERRA time series, using inverse FFT with interpolated and scaled magnitudes from 3) and appropriate phase angles. The algorithm for finding phase angles is described in Section 3.4.2.
5. Transform y_{sim} to $M_{HF,sim}$, using predicted VI. Nearest neighbour, random forest and GB models were trained and used for prediction. For the former, only calendar month, hour of the day and time series extracted from the MERRA model (e.g. the VI of M_{HF}) were used as predictors. For the latter two, EOF time series for several meteorological variables were optionally used in addition.
6. The time series for the improved model is achieved as $M_{LF} + M_{HF,sim}$.

The volatility of the high-frequency component varies in time. By simulating a stationary time series y_{sim} which was subsequently transformed to $M_{HF,sim}$ using the predicted VI, the varying volatility of measurements could be reproduced in the simulations. If was found that a GB machine learning model was somewhat better than k nearest neighbour and random forests for predicting the VI. By using EOF time series as predictors in addition to calendar month, hour of the day and time series extracted from the MERRA model, the correlation between VIs of simulations and measurements was increased slightly. A scatter plot of observed versus predicted VIs for the GB model with EOFs as predictors is shown in Fig. 4.4.

Fig. 4.5 gives some results from the simulation. The histogram in Fig. 4.5a clearly shows that the magnitudes of the high-frequency fluctuations from the MERRA model are too small. For the improved model, however, these are in

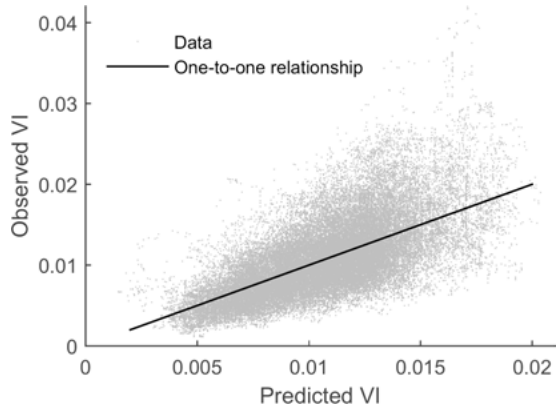

Figure 4.4. Scatter plot of observed vs. predicted variability indices (VI, a metric of the volatility of the high-frequency time series). The correlation is 0.67.

good agreement with measurements. Fig. 4.5b gives a representative example of the varying magnitudes of high-frequency fluctuations. The improved model can relatively well capture the magnitudes, although the match is not perfect. Figs. 4.5c–d show time series excerpts and PSD estimates for measurements, the MERRA model and the improved model. The fluctuation magnitudes from the improved model are more similar to those seen in measurements.

The underestimation of one hour step changes was the primary reason to improve the basic model described in Section 4.1.1. As Fig. 4.6 shows, the distributions of step changes for measurements and the improved model are very similar. Table 4.2, finally, gives RMS errors, correlations of VIs and some metrics describing one hour step changes. Overall, the improved model is really an improvement. Due to the filtering and addition of stochastic noise, the RMS error is slightly increased. The observant reader may have noted that the metrics given for the MERRA model are not identical to those presented in Section 4.1.1. The reason for this is that all six years of output (2007–2012) from the MERRA model was used for training, validating and testing the improved model.

4.1.3 Intra-hourly fluctuations

In many power system studies, wind data with a higher resolution than hourly is preferable [18, 111]. Two independent studies [112, 113] of the Irish power system compared the results from using hourly and sub-hourly (5–15 minutes) data in the simulations. Markedly higher levels of generator cycling and ramping rates were seen in the highly resolved simulations. Furthermore, hourly resolution models seem to underestimate the value of flexible generation and energy storage.

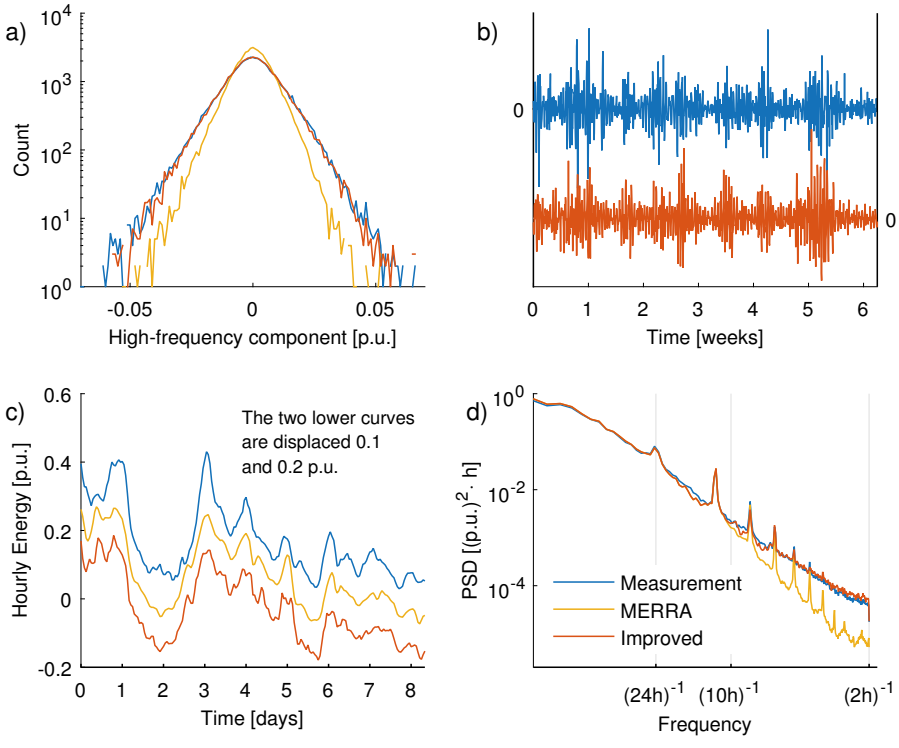

Figure 4.5. Comparisons of measurements and outputs from the basic (MERRA) and improved models. **(a)** Histogram of high-frequency components ($f > (10h)^{-1}$) using bin width 0.001 p.u., **(b)** representative example (in terms of correlation of variability indices) of high-frequency components, **(c)** example of hourly energy time series and **(d)** power spectral density estimates.

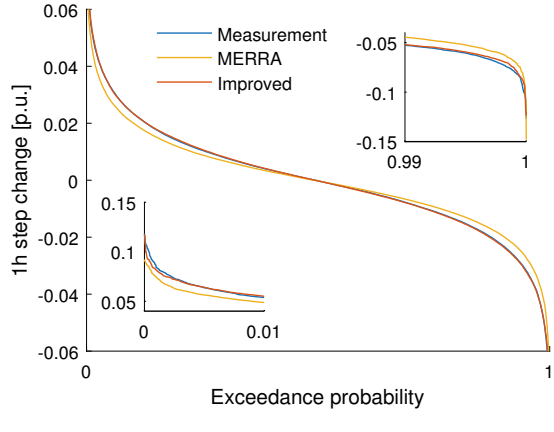

Figure 4.6. Duration curves for one hour step changes of the hourly generation time series (with the tails as insets). With the improved model, the step changes are no longer understimated.

Table 4.2. *Various metrics comparing the performance of the original MERRA model and the improved models. ΔP_{1h} is short for one hour step change in hourly energy and SD for standard deviation. All metrics except correlation are given in relation to the installed capacity, i.e. 3.6% corresponds to 0.036 p.u.*

	Measure-ment	MERRA model	Improved model (GB)	Improved model (GB, no EOFs)
RMS error	-	3.6%	3.8%	3.8%
Correlation VI	-	0.58	0.67	0.65
SD ΔP_{1h}	2.0%	1.7%	2.0%	2.0%
1^{st} %-ile ΔP_{1h}	-5.3%	-4.5%	-5.2%	-5.2%
99^{th} %-ile ΔP_{1h}	5.4%	4.9%	5.5%	5.5%

For many power systems, only hourly measurements are available. The meteorological models described in Section 2.2 are also often of hourly or poorer resolution. In Paper III, a methodology for simulating intra-hourly fluctuations (deviations from hourly mean values) was therefore developed. The ideas were relatively similar to those used for simulating hourly fluctuations described in Section 4.1.2; isolate the high-frequency part of actual measurements, simulate similar (stochastic) fluctuations using the IFFT algorithm with a search for phase angles, taking into account the non-stationarity of the data (right volatility at the right time) and finally add the simulated noise to smoothened data for the objective power system. Since no intra-hourly measurements were available for the Swedish power system, models were trained and evaluated using data from Denmark and Germany (5- and 15-minute temporal resolution respectively). The PSDs of these systems are similar to that for Sweden for frequencies up to $(2h)^{-1}$ and it can therefore be expected that the PSDs are also similar for higher frequencies.

As compared to the methodology described in Section 4.1.2, the approach differed in some aspects. Firstly, no EOFs were used for prediction of volatility since the EOFs are system specific. Only hour of the day, calendar month and variables derived from the hourly energy production time series were used as predictors. Secondly, instead of predicting the variability index, all observations were sorted into bins depending on the predictor values. Distributions (t location-scale) were fitted for each bin of the training data. The normally distributed and stationary simulated time series were transformed into simulated high-frequency components using the inverse of these transformations. Finally, a spline adjustment (see Section 3.3.3) was used to get correct hourly mean generation. Models trained with data from Denmark and Germany were used to simulate time series for both these systems. For our purposes, i.e. using these models for simulating Swedish intra-hourly fluctuations, the most interesting results are for the Danish model evaluated on the German system and vice versa.

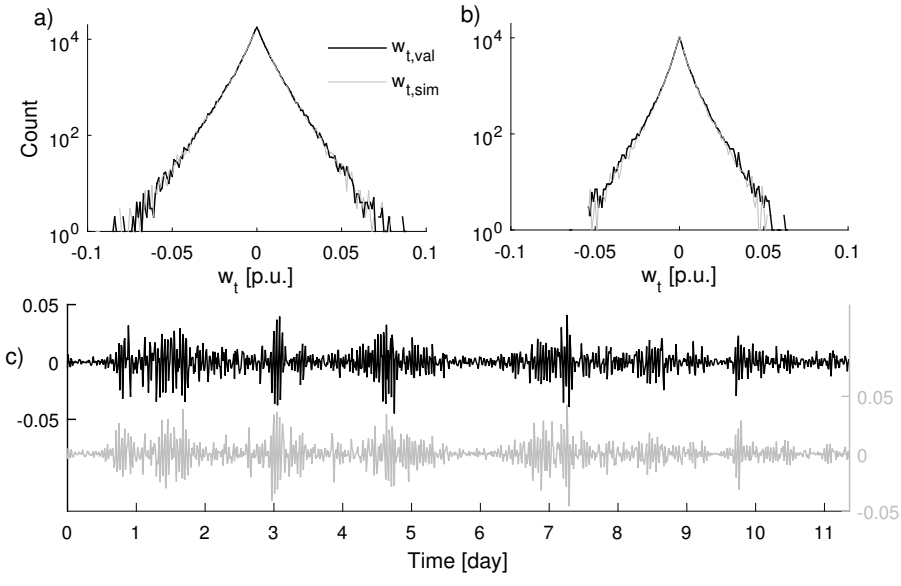

Figure 4.7. Similarity of deviations from hourly means (w_t) for simulations and validation data. **(a)** Histogram with Denmark as both training and objective system (Denmark-Denmark), **(b)** histogram (Denmark-Germany) **(c)** representative time series example (Denmark-Germany).

Let us begin by having a look on the simulated and measured deviations from hourly mean values ($w_{t,sim}$ and $w_{t,val}$). Fig. 4.7a–b give histograms of w_t for Denmark as both training and objective system (however using different time periods) and for Denmark as training system and Germany as the objective system. These combinations were chosen based on their, as an overall judgement, best and poorest performance respectively. In both cases, the simulated and measured deviations have very similar distributions, also in the tails. Fig. 4.7c shows a time series example of $w_{t,sim}$ and $w_{t,val}$ for Denmark as training system and Germany as the objective system. When comparing to the results in Section 4.1.2, the varying volatility of w_t for simulation and measurements are much higher correlated (0.94–0.95 as compared to 0.65–0.67). One reason for this is that $w_{t,val}$ is strongly dependant on the slope of the hourly energy generation P_h. When P_h is increasing, $w_{t,val}$ is likely to be strongly negative in the beginning of the hour and strongly positive in the end of the hour. When P_h is nearly constant, $w_{t,val}$ is often small in magnitude. This behaviour is easy to mimic in the simulations.

Fig. 4.8a shows a comparison between the simulated output $P_{t,sim}$ and measurements $P_{t,val}$. By visual inspections these time series have similar characteristics. The PSDs in Fig 4.8b–c strengthen the impression that the simulations were successful. As can be seen in Table II in Paper III, the SDs of one sample step changes are also very similar for $P_{t,sim}$ and $P_{t,val}$.

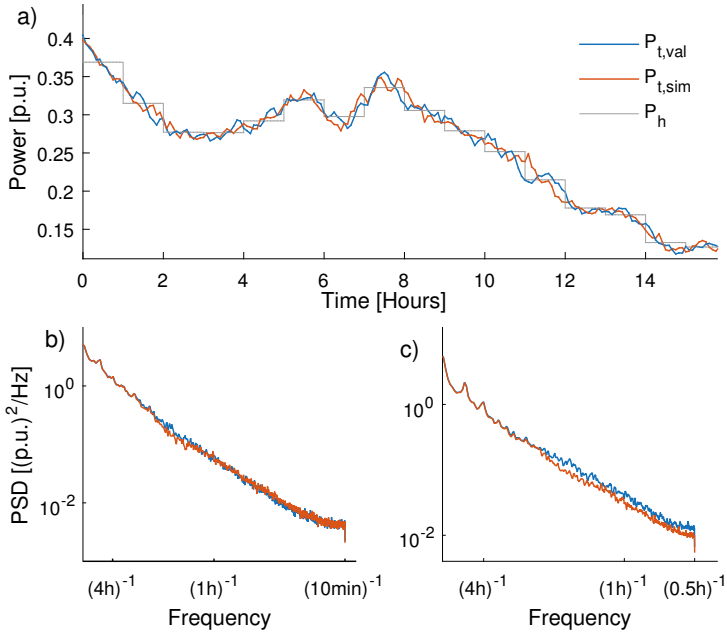

Figure 4.8. **(a)** Time series example of the simulated output $P_{t,sim}$ and measurements $P_{t,val}$ (Denmark as both training and objective system). **(b–c)** Power spectral density estimates for **(b)** Denmark-Denmark and **(c)** Denmark-Germany as training and objective systems respectively.

4.1.4 Synthetic forecasts

In Paper IV, a new method was developed for creating synthetic wind power forecasts for wind integration studies, see also Section 3.2.5. The idea is to use reforecasts from the GEFS dataset and some statistical processing for producing the synthetic forecasts. The modelled FEs can subsequently be computed as

$$\textit{Modelled FE} = \textit{Modelled generation} - \textit{Synthetic forecast}, \tag{4.1}$$

where the modelled generation is based on MERRA data as described in Section 3.2.1. Before evaluating the similarity of modelled and actual FEs, let us have a look at the terrain dependence and temporal consistency of modelled FEs.

Fig. 4.9a–b show how the SDs of D+1 and D+5 FEs vary over Sweden. Results are given for generation calculated from "raw" GEFS wind speed forecasts, i.e. before the addition of multivariate ARMA noise and removal of a fixed share of the error each hour. When comparing to the complexity of the terrain, quantified in Fig. 4.9c as the SD of ground heights within each grid cell, it is obvious that the day-ahead (D+1) errors are strongly correlated to the terrain complexity. FEs for longer horizons are, in contrast, not so terrain dependent, but are larger for offshore farms due to higher CFs[1]. Although we did not have enough measurements to properly evaluate if the relative differences in FE magnitudes were correct, the results seem reasonable when comparing to an earlier study [88] and FEs for three farms in Sweden.

An advantage with reforecasts, i.e. forecasts produced from historical data using approximately consistent models and assimilation systems, as compared to historical, operational forecasts is that the quality does not change much over time. In Paper IV, 30 years long generation and forecast time series were computed for 50 random locations in Sweden. In general, monthly CFs agree well except for year 2010–2012, which is likely due to a change in source for GEFS initial conditions. The quality of the forecasts only improves slightly over the years. A linear fit to SDs of hourly FEs computed for each year gives a reduction from 0.179 p.u. year 1985 to 0.161 p.u. year 2014 (average for the individual farms). The same analysis performed for aggregated forecasts gives a reduction from 0.078 p.u. to 0.075 p.u.

Let us now turn to a comparison between modelled and actual FEs for the whole of Belgium, see Fig. 4.10. An example of modelled generation and synthetic D+1 and D+7 forecasts is first presented in Fig. 4.10a. Note that towards the end of day two, there is a phase error in the D+1 forecast; the generation ramps up slightly earlier than the forecast. This type of errors, also present for operational forecasts, are generally not captured with purely

[1]When the errors are expressed in relation to the energy production, and not in relation to installed capacity as in Fig. 4.9, the opposite however holds.

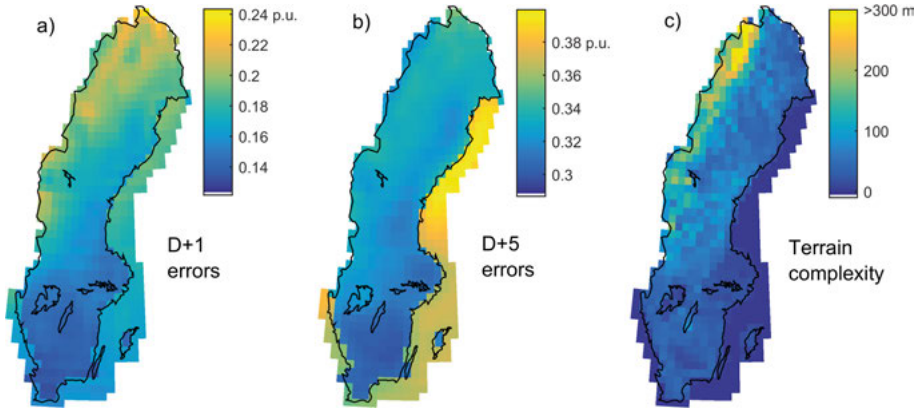

Figure 4.9. **(a–b)** Standard deviations of forecast errors for D+1 and D+5 respectively ("raw" forecasts). **(c)** Terrain complexity quantified by the standard deviation of ground heights within each grid cell. The ground heights were obtained from a database with 0.5 km horizontal resolution.

statistical methods for generating synthetic forecasts. Also note the abrupt change in the D+7 forecast at the beginning of day one. Similar step changes can be seen when actual forecasts are updated, especially for longer horizons.

Fig. 4.10b shows how the FEs for operational and synthetic forecasts depend on horizon. Note that only D+1 and the D+7 forecasts were available from the Belgian TSO. The good agreement should not come as a surprise since the model parameters were tuned to give a good match. In Fig. 4.10c, the impact from generation level on FE magnitude is displayed. For D+1, model and measurements both give highest errors when the national generation is around 0.2–0.7 p.u. During these occasions it is likely that many WTs are operating in the steepest part of the PC. For D+7, the curves are not agreeing as well, but both model and measurements give lower errors for very low national generation levels. Quantile-quantile plots for measured and modelled FEs are given in Fig. 4.10d–e. The distributions agree well, also in the tails. This is beneficial since the extreme errors are important for dimensioning of the reserves. The D+1 autocorrelations given in Fig. 4.10f have an excellent match. For D+7, the autocorrelation is however somewhat overestimated by the model. Fig. 4.10f, finally, shows PSD estimates for model and measurements.

For Sweden, day-ahead forecasts were available for three individual farms. The same ARMA model was used to generate noise for the synthetic forecasts, but slightly more of the error for each hour was subsequently removed (higher e_{fix} parameter). Although somewhat better performance (i.e. similarity to measurements) could be obtained by using separate ARMA models, we believed it to be beneficial to show that one model works well for both countries and both for individual farms and for national output, thereby demonstrat-

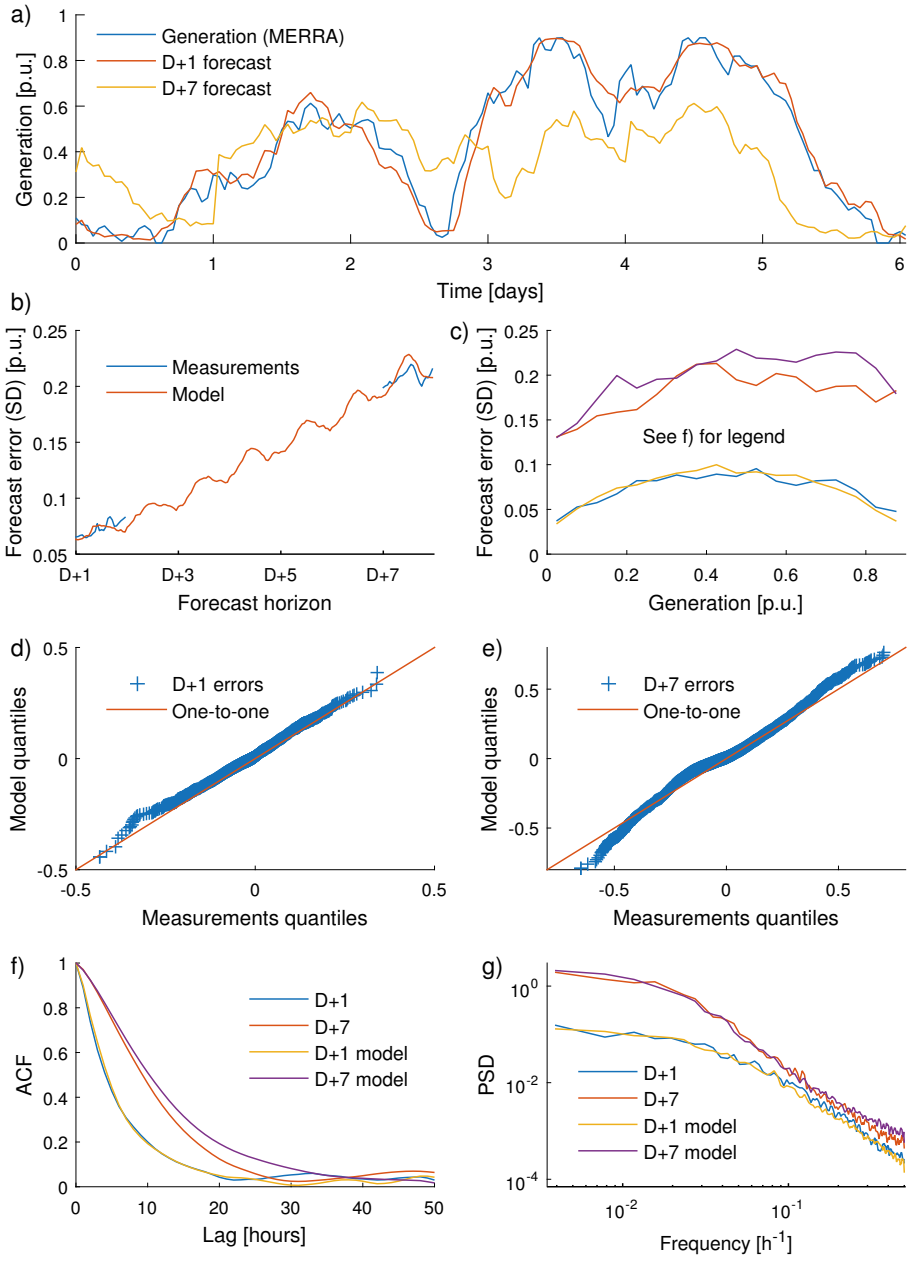

Figure 4.10. Evaluation of model performance for Belgium. **(a)** Time series example of modelled generation and synthetic D+1 and D+7 forecasts. Sub-figures **(b)**–**(f)** compare different properties of measured and modelled forecast errors (FEs). **(b)** Standard deviation (SD) of FEs depending on horizon (measurements are only available for D+1 and D+7), **(c)** Errors depending on generation level, **(d–e)** Q-Q plots, **(f)** Autocorrelation and **(g)** Power spectral density.

Table 4.3. *Impact from WT characteristics on modelled forecast errors. Results for hypothetical farms in Sweden for horizon D+1 and D+7. S1 is the baseline scenario, S2 has lower specific rating, S3 has higher assumed wind speed and S4 has 30% offshore capacity. All results are given as standard deviations of forecast errors, normalised to annual energy yield.*

	D+1	D+7
S1	0.193	0.510
S2	0.177	0.465
S3	0.177	0.466
S4	0.158	0.448

ing the robustness of the method. The most important result for Sweden was that when e_{fix} was tuned to give a correct magnitude of the errors for all farms combined, the magnitude of the errors for individual farms are also appropriate. Farm 1 is located in complex terrain and has the highest errors. Farm 2 is an offshore farm with higher CF than the other two farms and relatively high FEs. Farm 3, which has lowest errors both for measurements and model, is an onshore farm in somewhat hilly terrain, bordering flat terrain near the sea in the prevailing wind direction.

In Paper D, it was shown that the variability of wind power generation can be reduced significantly by increasing the average CF, which can be accomplished by, for example, increasing the share of offshore farms, increasing the hub heights or lowering the specific ratings. In Paper IV we investigated whether the reduced variability also leads to lower FEs. Four different scenarios were considered: S1 with a CF around 0.30 and S2–S4 with CFs around 0.36. S1–S3 consist of 50 randomly selected onshore farms in Sweden. The same farms were used for all these scenarios. In S4, fifteen of the farms were replaced with five offshore farms. As compared to the base scenario (S1), the higher CFs in S2–S4 were obtained by reducing the specific rating (S2), increasing the hub heights and thereby the mean wind speeds (S3) and by assuming 30% of the installed capacity as offshore farms (S4).

As can be seen in Table 4.3, increasing the CF not only leads to lower variability, but also to lower FEs[2]. The lowest errors are obtained for S4, which can be explained by longer average separation distance between the farms and different weather patterns onshore and offshore. This leads to lower correlation between the FEs of individual farms and consequently lower aggregated errors. The smoothing effect is strongest for shorter horizons when FE correlations are lower.

[2]Both these conclusions are valid when the results are normalised to the annual energy yield which, in my opinion, is the preferable option when comparing different scenarios.

4.2 Wind power in Sweden

In Section 4.2.1, results are given for historical wind power generation and trends of wind farm characteristics in Sweden. The analysis is based on hourly, aggregated measurements from 2007 to 2014, see Section 2.1, and a database of all operating wind farms. In Section 4.2.2, a selection of results from modelling future wind power generation in Sweden are presented.

4.2.1 Historical generation

Beginning with trends[3], Fig. 4.11 displays how hub heights, specific rating and CFs have evolved over time. The increase in hub heights has been more or less linear, from below 40 m in 1990 to over 100 m today. Modern WTs also have smaller installed capacity in relation to rotor areas; slightly above 300 W/m^2 as compared to around 400 W/m^2 in the 90s and early 00s. The development is not linear as for hub heights, but rather it seem to be a change of trend around year 2005. This is likely related to the increased deployment of farms in forests at the same time. Increased hub heights, relatively larger rotors and, to a smaller extent, improvements in WT efficiencies have led to markedly higher CFs. When a farm is deployed, the estimated annual energy yield is reported to the electricity certificate system [60]. Figure 4.11c shows how the estimated CFs have increased from around 0.20 to 0.33. As shown in Figure 4.11d, the measured national CF has increased from around 0.21 year 2007 to around 0.29 year 2013–2014 (long-term corrected data). The measured CF has increased faster than expected from the estimated CFs which can be explained by the systematic overestimation of energy yields for farms deployed in the 90s and early 00s. In year 2014, these old WTs only constituted a small fraction of the total capacity and consequently the overestimation of the national CF was much smaller than in 2007.

Figure 4.12 and Table 4.4 give information on the distribution of hourly generation. The generation very seldom, on average 14 hours per year, exceeds 0.8 p.u. The 90$^{\text{th}}$ percentile is only 0.51 p.u. If high wind penetration is an issue, it is therefore not too costly to curtail wind power at times with low load and high wind (especially since the price can then be expected to be low). Some quantitative examples are given in Section 4.2.2.

There are significant seasonal differences in wind generation. Diurnal patterns are also present, although not as strong as the seasonal variations. Fig. 4.13 illustrates the patterns. In Fig. 4.13b, the monthly means have been removed in order to visualise the diurnal cycles each month more clearly. These pattern, which are not properly captured by the MERRA reanalysis, seem to be related to the rise and set of the sun.

The variability of intermittent energy sources is not a problem per se; if the fluctuations are correlated to those of the electric load it can actually be

[3]The values are weighted on the installed capacity of each farm.

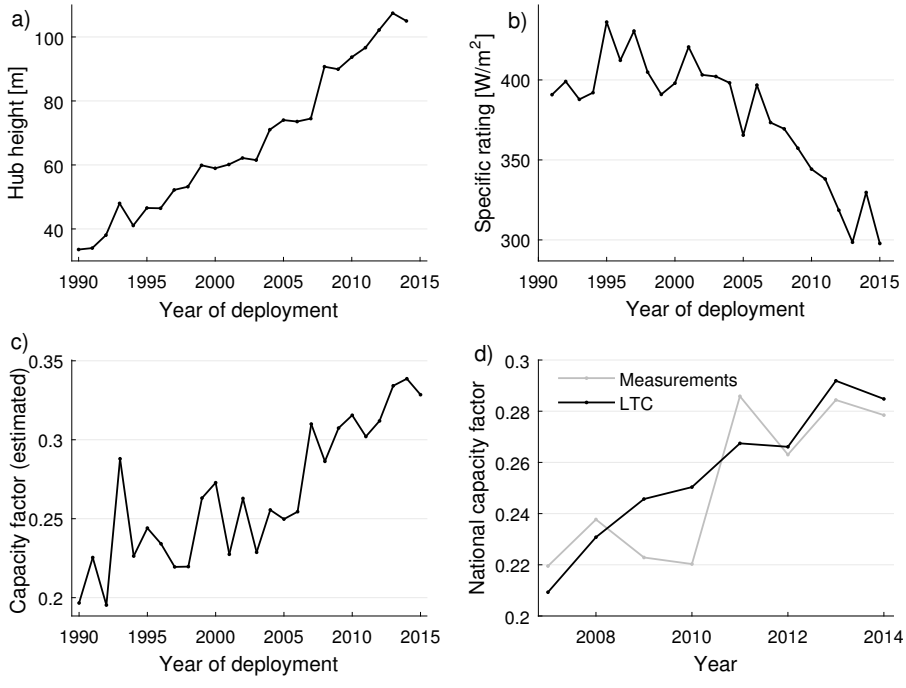

Figure 4.11. Capacity weighted trends for Swedish wind power, see Paper D for more information on the dataset. **(a)** Hub heights, **(b)** specific ratings, **(c)** capacity factors (estimated by farm owners before construction) and **(d)** measured national capacity factors, both raw and long-term corrected (LTC) data.

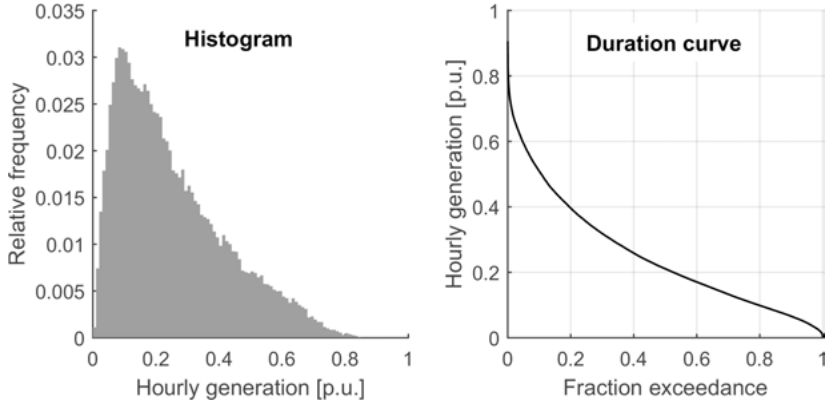

Figure 4.12. Figures illustrating the distribution of Swedish hourly wind generation 2007–2014. **(a)** Histogram computed with bin width 0.01 p.u. and **(b)** duration curve.

Table 4.4. *Distribution metrics for wind power in Sweden (hourly measurements 2007–2014). The time series has been normalised to the installed capacity at each time step. The 10th percentile is considered "firm capacity" by the Swedish TSO.*

Metric	Value
Capacity factor	0.25
Standard deviation	0.17 p.u.
Minimum	0.003 p.u.
1st percentile	0.021 p.u.
10th percentile	0.066 p.u.
90th percentile	0.51 p.u.
99th percentile	0.71 p.u.
Maximum	0.90 p.u. (some years below 0.8 p.u.)

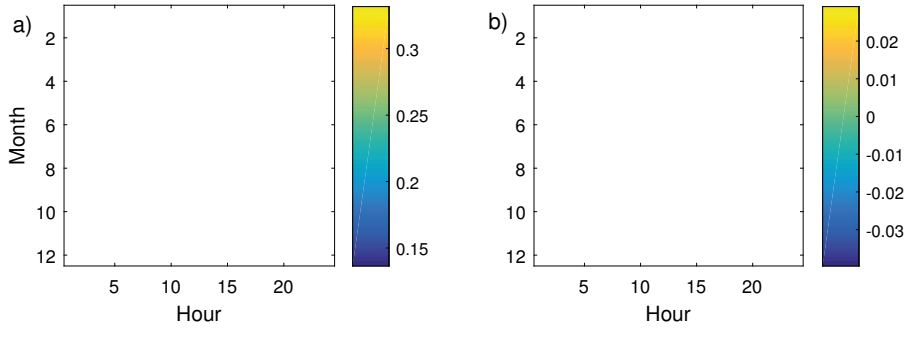

Figure 4.13. Seasonal/diurnal plots for wind power in Sweden (2007–2014). **(a)** Average capacity factors (CFs). The generation is considerably higher in wintertime. **(b)** Deviations from monthly CFs. The pattern seems to be related to the rise and set of the sun.

advantageous. As can be seen in Fig. 4.14a, the seasonal variations of wind power and load in Sweden are in good agreement. In the figure, both wind generation and load have been normalised to their average values. Unfortunately, a closer look at the wind-load dependence reveals that at times with very high load, the wind generation is below its average. In Fig. 4.14b, the mean and 10^{th} percentile of wind generation are plotted as functions of the load (grouped into 100 bins). As an example, the first load bin correspond to the 1% of the time when the load is lowest (around 9 GW). The mean wind generation is then around 70% of its average. Note that for the computation of 10^{th} percentiles of wind generation, five load bins at the time were considered in order to get more robust results. The 10^{th} percentile is an important metric since this is defined as 'firm capacity" by the Swedish TSO. Recently, the TSO's estimated firm capacity of wind power was increased based on observations in winter-time (when the load is generally higher). According to my analysis, this was a mistake; during the 5% of the time when the load is at its maximum, the 10^{th} percentile of wind power is actually very slightly below the 10^{th} percentile for the whole time period. A physical mechanism for this is suggested in [114]: *"...load extremes are often due to relatively infrequent large-scale high-pressure weather systems that typically bring calm winds"*. Indeed, very high loads in Sweden generally coincides with high pressures (national average pressure data from the MERRA reanalysis).

Fig. 4.14b illustrates the importance of using synchronous wind and load data. If one insist of using purely statistical methods for producing wind power time series, the relationship between load and wind generation must be taken into account. The following example quantifies the error from neglecting the relationship (but taking into account the seasonality). When randomly permuting[4] the years of the wind power time series, the generation was in average 20% above mean for the 5% of time with largest loads. This can be explained by the higher generation in wintertime. With synchronous data, the generation is 7% below average.

The variability of Swedish wind power is relatively low in comparison to that in other systems [6]. The main explanation for the low variability is the large mean distance between farms. In Table 4.5, a few metrics of change (one and four hour step changes) are given. The distributions of step changes are given in Fig. 4.3 on page 66. Kurtosis is one way to quantify the occurrence of relatively extreme values. Both one and four hour step changes have kurtosis around five, i.e. the distributions are significantly more heavy-tailed than normal distributions.

The average revenue for wind farm owners does not necessarily equals the time average of electricity prices. Firstly, wind energy may be produced at

[4]One thousand permutations were performed. Only the 382 permutations with no years in common were considered. As an example, [2009, 2008, 2012, 2013, 2007, 2011, 2014, 2010] would not be accepted since year 2008 is the second element both for the load data and for the permuted wind power data.

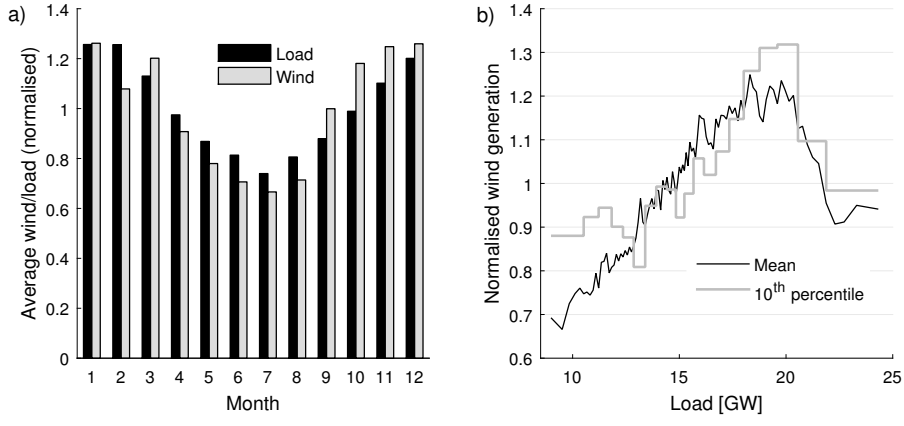

Figure 4.14. Relationships between wind power generation and load in Sweden. The data have been normalised to the means for the whole time period (2007–2014). **(a)** Wind has a similar seasonality as the load. **(b)** Wind generation for 100 bins of load. Each bin has the same amount of data, e.g. the first bin contains wind generation corresponding to the lowest percent of load observations. The 10^{th} percentiles of wind generation are given for five load bins at the time.

Table 4.5. *Metrics for one and four hour step changes for wind power in Sweden (hourly measurements 2007–2014).*

Metric	Value 1h	Value 4h
Standard deviation [p.u.]	0.019	0.061
Minimum [p.u.]	-0.13	-0.33
1^{st} percentile [p.u.]	-0.051	-0.16
99^{th} percentile [p.u.]	0.053	0.17
Maximum [p.u.]	0.12	0.41
Kurtosis	5.1	4.8

times with prices differing from the average, e.g. more wind energy is produced during winters. Secondly, if the penetration is significant, a high wind power generation at certain hours will have a negative impact on the price. The ratio of wind weighted price and time weighted price (i.e. a simple mean) is called value factor (VF). Wind value factors reported in the literature is often around unity for low penetrations but declines when more wind power is deployed. Hirth [7, 115] has showed that at a penetration level of 30%, wind electricity can be worth 10–50% less than generation from a constant source. The VF is system specific, e.g. depending on the relationship between wind and load and the type of generators used for balancing wind power.

According to data from the Swedish TSO [58], the wind power generation and load for year 2015 were 17 and 135 TWh respectively. Based on market data from NordPool [116], the average prices weighted on time, load and wind generation were 22.0, 23.3 and 21.3 €/MWh respectively. The VF for wind power is thus only slightly below unity. A slower decline in the VF with penetration rate can be expected for hydro dominated systems [115]. It can also be noted that in 2015, prices were low in a historical context and differed little between the bidding zones. An interesting topic for future studies (see Chapter 6) is how the VF in the Nordic system will change in a future with considerably higher wind penetration and furthermore if a potential VF reduction can be alleviated with proper measures.

4.2.2 Future generation

In this section, some results from modelling future wind power generation in Sweden are presented. The results are given for scenarios developed in Paper D and sometimes, for comparison, for linearly scaled historical generation data. The database of potential farms, used for generating the scenarios, is illustrated in Fig. 4.15. When comparing to the situation today (Fig. 4.15a), the distributions of onshore farms with all permits (Fig. 4.15b) and farms in the permitting process (Fig. 4.15c) are strikingly similar; most in bidding zones SE2 and SE3, somewhat less in SE4 and the least in SE1. A large share of the planned offshore farms are however located in SE4.

Two of the main scenarios from Paper D, A1 with 20 TWh/a and C1 with 50 TWh/a, are shown in Fig. 4.16. The A1 scenario, expected to be realised around 2020, has a very small offshore capacity but for C1, the offshore energy share is 26%.

As discussed in Section 1.1, duration curves of the net load can provide insights into the challenges of integrating variable renewables such as wind power. Fig. 4.17 shows duration curves for both hourly generation and one hour step changes for the load and the net load with 50 TWh/a of wind power (scenario C1). The mean net load is reduced by 5.6 GW, the maximum net load by 2.3 GW and the minimum by 9.1 GW. Consequently, if wind power

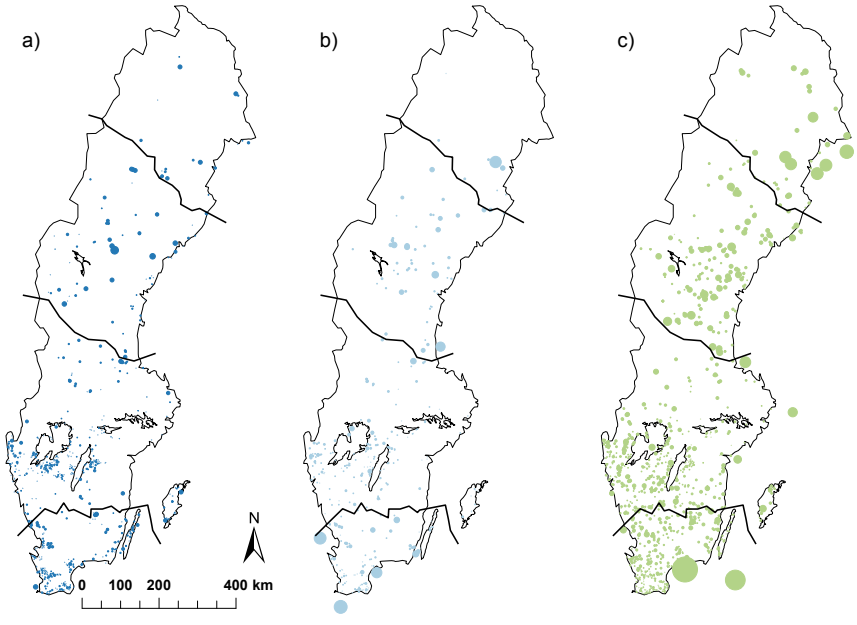

Figure 4.15. Database of wind farms in Sweden from Paper D. Bidding zone borders are indicated by thick lines (SE1 in north to SE4 in south). The circle areas are proportional to capacity. Farms **(a)** in operation or under construction, **(b)** with all permits in place, **(c)** in the permitting process.

replaces nuclear generation, some additional peak generating capacity might be necessary. The distribution of net load step changes also differ from that of the load, in particular for the tails. Ramping capabilities is however not the main issue for the hydro-dominated Nordic power system.

In order to alleviate the impact from wind power on the power system it is important to pinpoint the most influential parameters for the resulting variability and quantify costs and benefits of different strategies to reduce the variability. Some future studies on this topic are planned, see Chapter 6. In Fig. 4.18, results from a parameter study in Paper D are given. As a starting point, a scenario with 20 TWh/a onshore farms, average CF of 0.34 and a specific rating of 280 W/m^2 was used. The SD of one hour step changes was used as a proxy for variability. All results were normalised to the reference case.

According to Fig. 4.18a, adding more farms of the same type has a relatively small impact on the variability. The geographical smoothing effect is in other words small when the number of farms is significant to start with and no action is taken to steer the deployment of farms towards regions with little capacity. However, as can be seen in Fig. 4.18d, if the placement of farms is optimised, a considerably lower variability can be obtained with only a few farms. Depending on whether only available onshore projects, all onshore locations or onshore and near offshore locations are allowed in the optimisation,

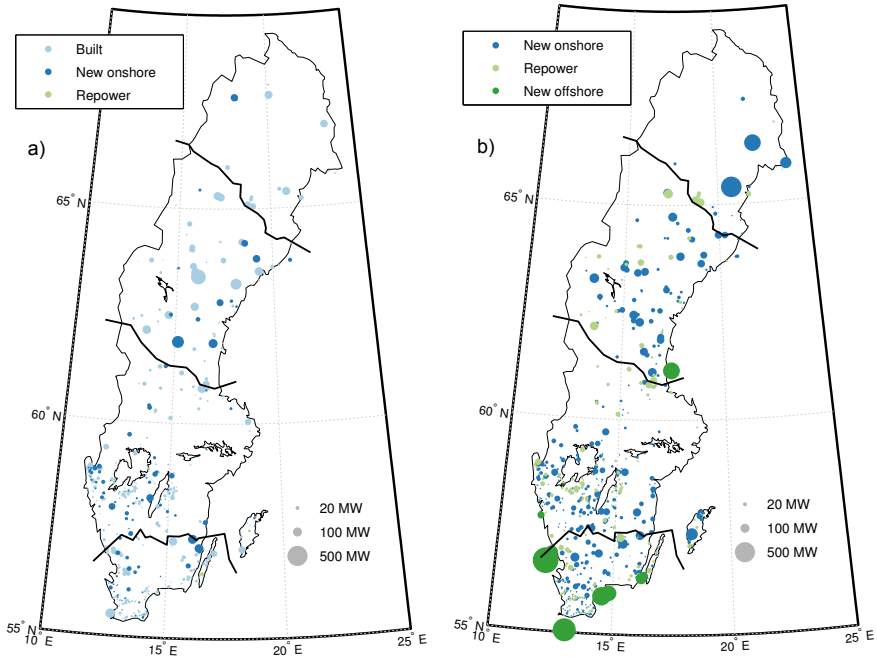

Figure 4.16. Two of the main scenarios from Paper D. "Built" indicates farms already in operation or under construction, "Repower" indicates current wind turbines (WTs) that have been replaced by larger, modern WTs. Bidding zone borders are indicated by thick lines. **(a)** Scenario A1: around year 2020, 20 TWh/a, 7.5 GW, **(b)** scenario C1: 50 TWh/a, 14.2 GW.

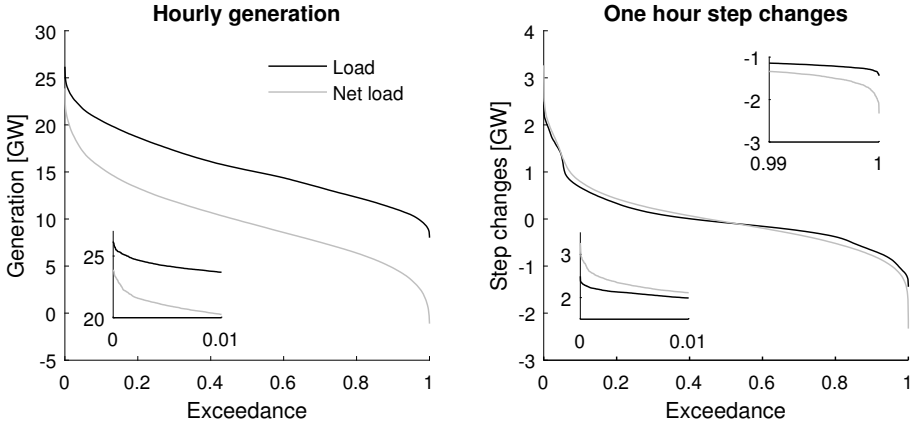

Figure 4.17. Duration curves of **(a)** hourly generation and **(a)** one hour step changes for load and net load. The results are based on measured load for 2007–2014 and wind power scenario C1 from Paper D. In this scenario, wind power supplies 50 TWh/a or around 36% of the load.

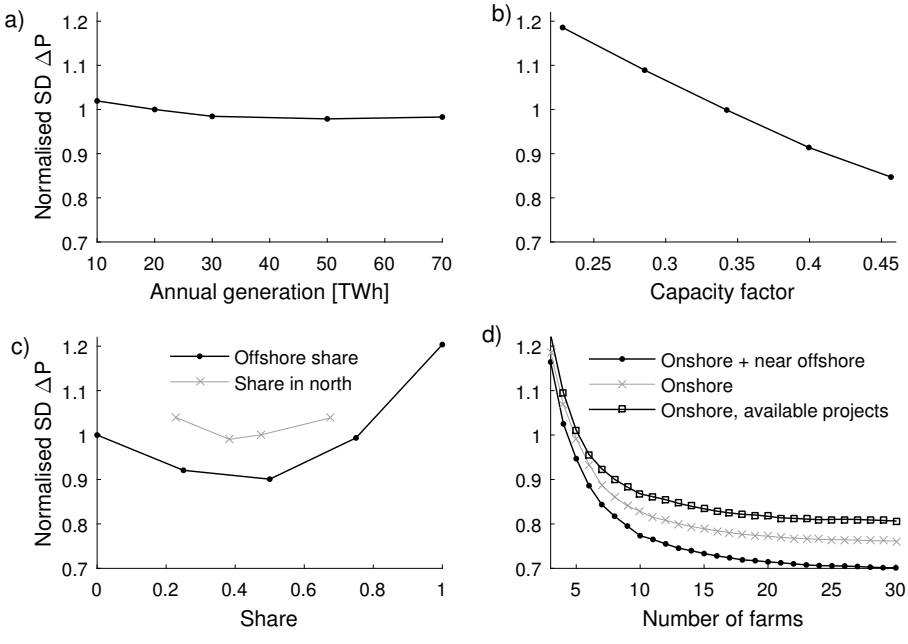

Figure 4.18. Impact from several parameters on the standard deviation of one hour step changes of wind generation. All results have been normalised to that for a reference scenario. **(a)** Adding more farms with the same characteristics has a small impact on variability, **(b)** variability is greatly reduced with higher average capacity factors, **(c)** minimum variability is attained with shares of offshore farms and farms in north around 40% and **(d)** optimising the geographical distribution of farms has the potential of significantly reducing the step change magnitudes.

the SDs of step changes saturates at around 70–80% of that for the reference case. Note that the same CF was assumed for all farms (including offshore) in order to isolate the effect of optimal geographical dispersion.

Fig. 4.18b shows the impact from the average CF on variability. Increasing the CF leads to a strong reduction of step change magnitudes. Fig. 4.18c quantifies the variability depending on share of farms in northern Sweden (SE1 and SE2) and offshore energy production respectively. The share of capacity in the north is not the most important factor, but lowest variability seem to be obtained with around 40%. Increasing the share of offshore wind power would reduce the variability. This can be explained both by an increase of the mean distance between farms and a higher average CF (0.46 was assumed for offshore farms). Too high shares of offshore wind is not desirable since the generation is then concentrated to only a few locations. A final note is that the specific rating has a very small impact on the step changes as long as the CF is not altered, i.e. from a variability point of view it is not important whether a certain CF is obtained by a low mean wind speed and a large rotor area or by a higher wind speed and a smaller rotor.

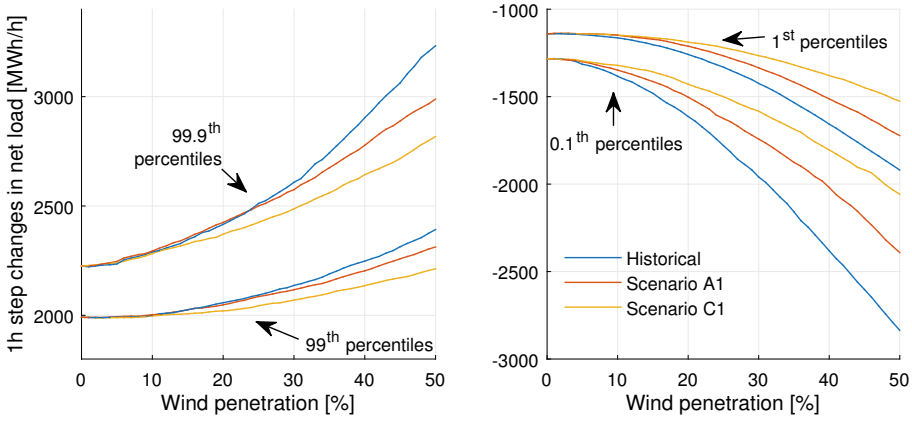

Figure 4.19. Impact from wind penetration and from using different wind power datasets on extremes of one hour step changes in the net load. The three wind power time series were linearly scaled. For comparison, the maximum load was 26 GW.

In Section 1.3 it was stated that linear scaling of historical wind generation might be very misleading for studying the future. The following exercise serves as an example. Let us begin by pointing out that the geographical smoothing for Sweden is close to saturated; building more farms of the same type that we already have in the same regions will reduce the normalised variability very little. It is thus reasonable to, as a first approximation, linearly scale generation in order to study different penetration levels of similar WTs. In Fig. 4.19, extreme values of net load step changes are given for different penetration levels. Three sources of wind data were used: historical measurements (as in [6]) and scenarios A1 and C1. The impact from using different datasets is quite dramatic, especially when considering low percentiles of step changes for high wind penetrations. Note that positive extreme ramps of the load are considerably larger in magnitude than negative ramps. The latter are however affected to a larger degree by increasing the wind penetration.

As in Section 1.3, it can be worth mentioning that the difference in step change magnitudes are not due to differences between modelling and measurements per se; modelling with the historical distribution and characteristics of WTs gives almost identical results as using measurements. The main lesson to be learned is that measurements should not be used for studying the future if one does not have good reasons to assume that CFs etc. will not change for future farms[5].

As shown in Table. 4.4 on page 79, it has historically been uncommon with generation above 0.7 p.u. and relatively uncommon with generation above 0.5 p.u. If a high instantaneous wind penetration is considered a problem, curtail-

[5]With that said, it is of course not easy to foresee the geographical distribution and characteristics of future farms. Scaling historical measurements implicitly assumes no change at all; most likely scenarios can do considerably better than that.

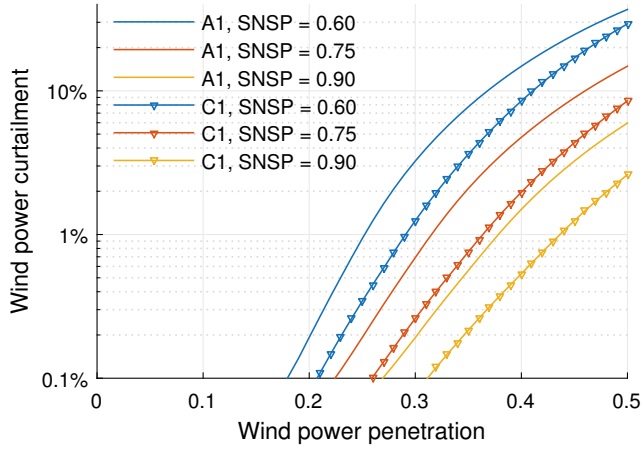

Figure 4.20. Curtailed wind energy for different average penetration levels (after curtailment). The amount of curtailed wind depends heavily on the allowed system non-synchronous penetration (SNSP) and the average capacity factor (0.31 for scenario A1 and 0.40 for C1).

ment of wind power is therefore a realistic and potentially an economically rational option. The system non-synchronous penetration (SNSP) is defined as [117]

$$SNSP = \frac{Wind - HVDC\ imports}{Demand - HVDC\ exports}. \tag{4.2}$$

In the following analysis, the curtailed wind energies necessary to keep the SNSP below certain values are calculated. The analysis was simplified by not considering HVDC links and by only studying Sweden. Measured loads for year 2007 to 2014 and wind power scaled from scenarios A1 and C1 from Paper D were used. For Ireland, allowed SNSPs in the range 60–75% are anticipated for year 2020 [117]. In other systems, e.g. Portugal and Denmark, higher instantaneous wind penetration have been observed [14]. The amount of curtailed wind energy for allowed SNSPs of 60%, 75% and 90% and wind penetration levels up to 50% are given in Fig. 4.20. Note that the penetration levels are after curtailment and that the curtailed energy is given in relation to the uncurtailed generation. For up to 20% average wind penetration, negligible amounts of wind would have to be curtailed, regardless of the allowed SNSP. With 50% penetration, the curtailed energy increases to high levels if not a SNSP of 90% is allowed. The average CF also has a relatively strong impact of the curtailed energy. Scenario A1 has a CF of 0.31, only slightly higher than today, while scenario C1 has a CF of 0.40.

Wind power curtailment deals with the issue of high wind and low load. We not turn to the opposite problem of system adequacy; what happens when the load is high and wind power generates little? Will future farms be better

Table 4.6. *Metrics for quantifying low wind / high load for historical wind power measurements as well as time series for two scenarios (A1 and C1). Firm capacity is the 10^{th} percentile in accordance with the Swedish TSO's guidelines. Firm capacities for high load is given for the top five percent of the load. The wind power time series were linearly scaled to 30 TWh/a or in average 3.4 GW.*

	Historical	A1	C1
Firm cap.	0.89 GW	1.0 GW	1.2 GW
Firm cap. high load	0.88 GW	1.0 GW	1.2 GW
Max load - max net load	1.0 GW	1.2 GW	1.5 GW

at handling such situations? The preferred method [114] for determining the "capacity value" of wind power is to compute the constant increase in load that gives the same loss of load expectation (LOLE) as before wind power was added. These calculations require information on accepted LOLE, forced outage rates for conventional generators and import capacities. Since these data were not available, three other metrics were computed:

1. Firm capacity, defined as the 10^{th} percentile in accordance with the Swedish TSO's guidelines.
2. Firm capacity for the 5% of the time with highest loads.
3. The difference between maximum load and maximum net load.

As before, the computations were performed for three different wind power time series: historical measurements and scenarios A1 and C1 from Paper D. The time series were linearly scaled to 30 TWh/a or a mean generation of 3.4 GW. The CFs were 0.25, 0.31 and 0.40 respectively. As can be seen in Table 4.6, the firm capacity is highest for scenario C1. Firm capacities are very similar for high loads and for all data. The differences between maximum loads and maximum net loads are around 20% higher than the firm capacities. The latter metric is very sensitive to the wind generation at hours with very high load. For individual years this metric varies between 0.8–3.2 GW for all datasets. This is consistent with the high variations in capacity values, see [114] for a few examples. For robust assessment of the contribution from wind power to system adequacy it is thus valuable with long time series, e.g. 35+ years from reanalysis data.

4.3 Variability

Besides the development of wind power models (Papers I–IV), one review article and two studies on variability/correlations are included in this thesis (Papers V–VII). Some results and conclusions from the review are given below. Results on wind power correlations in Europe are presented in Section 4.3.1.

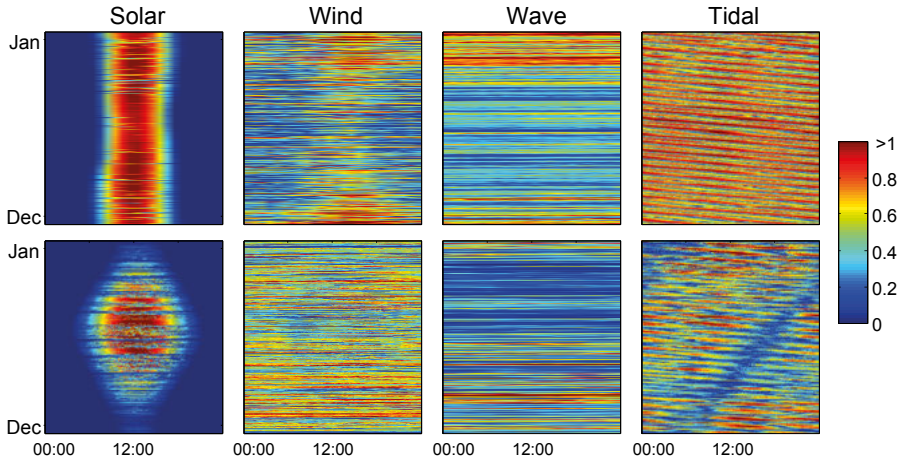

Figure 4.21. Site specific examples of variability at two sites per energy source (solar irradiation, wind speed, significant wave height and tidal current speed). The temporal resolutions differ from one minute to one hour. The measurements are normalised to the 98th percentile measured for each site.

Section 4.3.2, finally, gives results on net load variability in a highly or fully renewable Nordic power system.

As a background for future studies, a review on variability and forecastability of non-dispatchable renewable energy sources was conducted in collaboration with solar, wave and tidal power researchers (Paper V). The aim of the paper was to compare the methodologies used for the different sources. The foci were primarily on temporal variability and the effects of aggregation and not so much on spatial variability.

Although both wind and wave power is driven by energy from the sun, the temporal characteristics of these three sources are strikingly different. Solar irradiation has a clear maximum defined by geographical location and time. Cloud movements can introduce an almost binary pattern, where the point irradiation switches between almost zero and maximum. Because of this, research focus is primarily on short time scales. The wind speed can also have seasonal and diurnal patterns, although the strength of these vary from site to site. Stochastic variations of the wind are substantial and it has been shown that wind variations on the time scale of 1–6 hours can be challenging for the power system. Waves are induced by the wind, but the variations are smoothened out. The diurnal variations are therefore small. Tidal currents finally are driven by the gravity of the moon, and follows very regular patterns. Commonly, tides are semi-diurnal with two high tides and two low per day. Fig. 4.21 shows examples of variability of the different sources during one year.

Of great importance is also the correlation of production from spatially distributed sources. High correlations imply large variations for a distributed fleet of generation units, while lower correlations give a smoothing of the aggregated power. The correlation of wind power often exhibit an approximately exponential decay with separation distance. Averaging over longer time-periods gives higher correlation, while the short-term variations are more independent. The solar irradiance has a substantially higher correlation than wind power for similar separation distances.

Forecasting of intermittent renewables becomes increasingly important as the penetration levels increase. A lot of research and practical experience is available, not the least for wind power. Forecasting systems, often using physical NWP models and statistical post-processing, are operational in several countries, and the results have improved considerably over the past years. A conclusion from the review was that the accuracy of the forecasts are often hard to compare since different metrics are used for the different fields.

A general conclusion in Paper V was that more studies on combinations of the sources would be desirable. The disciplines could also learn from each other and benefit from the use of more unified methods and metrics.

In Papers VI and VII, generation time series were separated into frequency components. This gave us the possibility to study the variability characteristics in more depth. The same cut-off frequencies were used:

- Long-term component (T > 4 months)
- Mid-term component (2 weeks < T < 4 months)
- Mid/short-term component (2 days < T < 2 weeks)
- Short-term component (T < 2 days)

An example of a filtered wind power time series is shown in Fig 4.22. Appropriately filtered time series have two desirable properties: the sum of the frequency components add up to the input signal and the sum of the variances add up to the variance for the original (raw) time series (see Section 3.5). As an example, the SD of Swedish wind generation 2007–2014 was 0.17 p.u. The contributions were 0.054, 0.12, 0.083 and 0.073 p.u. from the short, mid/short, mid and long-term components[6]. Fluctuations with periods ranging from two days to two weeks are thus contributing most to the SD of the raw signal. The same analysis for one and four hour step changes shows that the short-term component is clearly dominating. If one is interested in more details on which frequencies contributes mostly to the variance, a cumulative PSD plot can be informative, see Section 3.4.1.

[6]Note that standard deviations (SD) does not sum in the same way as variances; the SD of the raw signal is, to a close approximation, the square root of the sum of the squared SDs of the constituent components

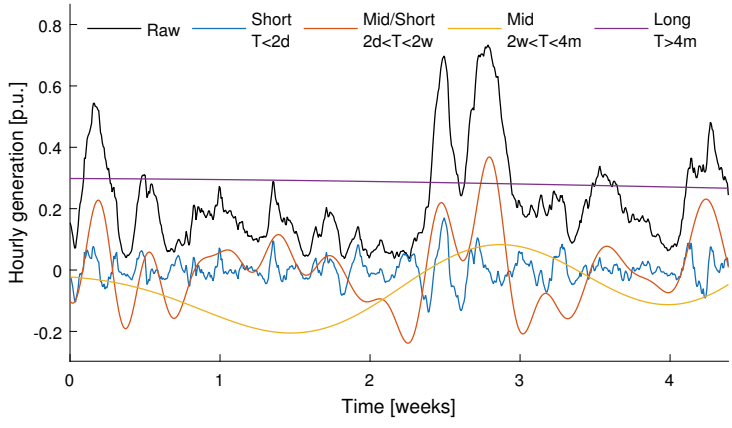

Figure 4.22. Filtering of a wind power time series. The raw signal is decomposed into four frequency components.

4.3.1 Wind power correlations

The correlation between wind power generation in different countries is important since it helps us quantify the reduction in aggregated variability when electrically interconnecting the countries. In Paper VI, hourly, country-wise time series of wind power output were generated for all European countries (or groups of smaller countries) and linear correlation coefficients ρ were studied. In order to deepen the analysis, ρ's were not only computed for these time series, but also for one hour step changes and for band-pass filtered data.

As compared to the methodology outlined in Sections 3.2.1 and 3.2.2, a simplified approach was taken to develop scenarios for year 2020 and model hourly output. The methodology is not given here, see Paper VI for details. Overall, the results were however (surprisingly) good when comparing to measurements and it was concluded that the simplified method is adequate for studying correlations. In particular, there is a good match in ρ's for hourly energy, step changes as well as for all band-pass filtered time series.

It is well known that the outputs from wind farms are correlated and that, in general, ρ decreases with longer separation distances [36]. Exponential models are commonly used to describe this relationship, see Eq. (3.8) on page 39. Most often, correlations between outputs of individual farms or wind speed measurements have been studied. The same analysis can be performed with nationally or regionally aggregated wind power time series [37, 41].

Pair-wise correlations versus distance and an exponential fit to the data are shown in Fig. 4.23a. Apparently, the exponential model is adequate (also) for country-wise generation. Fig. 4.23b shows smoothing spline fits of correlations versus distance for hourly energy, one hour step change and long, mid, mid/short and short-term components respectively. The correlations between the long-term components are generally high. This is expected since the

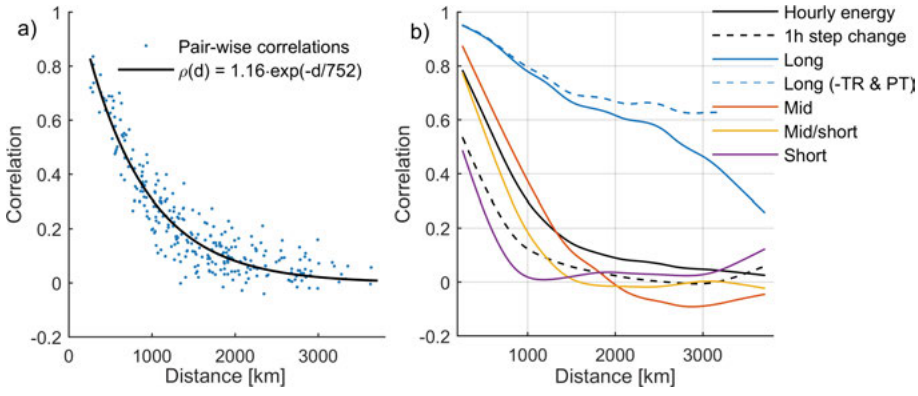

Figure 4.23. Correlations between wind power generation in the European countries depending on separation distance. **(a)** Pair-wise correlations and an exponential fit to the data (hourly energy). **(b)** Smoothing spline fits of correlations for hourly energy, one hour step change and long, mid, mid/short and short-term components respectively. The dashed blue line shows the correlations between long-term components when excluding Turkey and Portugal.

seasonal wind patterns in Europe are similar, with higher generation in wintertime. The spread in the results is however very large; ρ's for countries 2000–2500 km apart, for instance, vary between 0.17 and 0.83. The largest part of the spread can be attributed to the much lower correlations seen for Portugal and Turkey, which have distinctly different seasonal patterns than the rest of Europe. For fluctuations with shorter periods, ρ gradually decreases. In particular, the short-term variations are even less correlated than the step changes. In comparison with the results for individual wind farms (see e.g. [41]), ρ is markedly higher, especially for step changes.

For a single country, the correlations can be visualised in a map. Maps for all countries are provided in Paper VI. Correlation maps for Sweden are shown in Fig. 4.24.

4.3.2 A fully renewable Nordic power system

In Paper VII, the net load variability in the Nordic power system was studied. Different time scales and various shares of solar, wind, wave and tidal power were considered. The hourly net load was calculated from historical consumption for year 2010 through 2014 and meteorological resource data, using scenarios for likely deployment across the Nordic countries. For wind power, the methods described in Sections 3.2.1 and 3.2.2 were used to generate scenarios and model hourly generation.

The principal research question was how the net load variability will increase for a highly and fully renewable power system, assuming that new renewables replace fossil and nuclear generation on an equal energy basis. Fur-

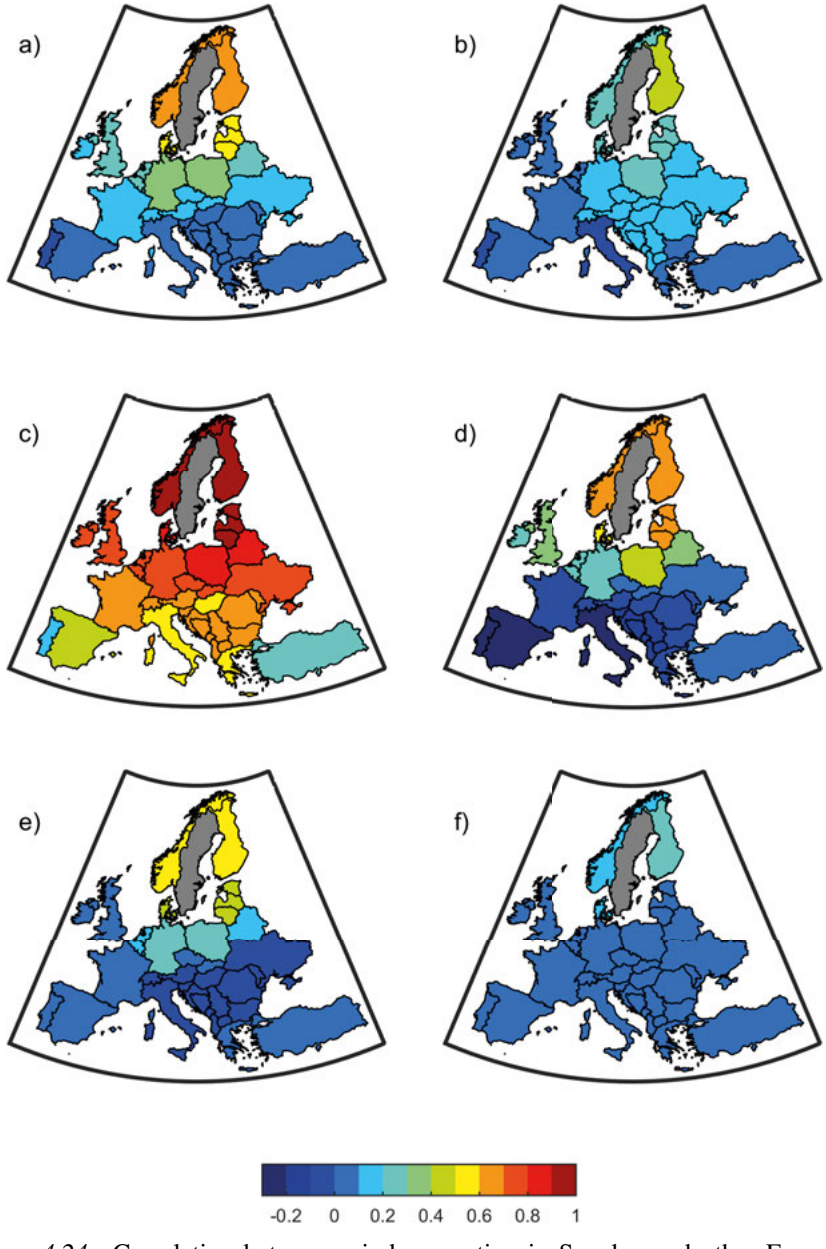

Figure 4.24. Correlation between wind generation in Sweden and other European countries. **(a)** Hourly energy, **(b)** one hour step change, **(c)** long-term component, **(d)** mid-term component, **(e)** mid/short-term component, **(f)** short-term component. The correlations are highest for the long-term components and lowest for the short-term components. Similar maps for all countries are provided in Paper VI.

thermore, it was studied if a wise combination of the sources can reduce the variability. As for the correlations described in Section 4.3.1, the analysis was performed for hourly energy, one hour step changes and the four frequency components specified on page 90. The net load (NL) was defined as

$$NL = load - IRE - thermal - nuclear = hydro + imports - exports, \quad (4.3)$$

where IRE is short for intermittent renewable energy[7]. In the highly renewable scenario, 20% of the annual energy was obtained from IRE sources. All fossil and 30% of the nuclear generation were assumed to be dismantled. In the fully renewable scenario, all fossil and nuclear generation was replaced with renewables, resulting in 36% energy from IRE sources.

As already seen in Fig. 4.21 on page 89, PV, wind, wave and tidal have different generation patterns. Fig. 4.25 shows the SDs of the four frequency components for the considered IRE sources and for the electric load. The distinctly different variability characteristics are now even more obvious. The overall variability is lowest for the load, relatively similar for wind, wave and tidal and largest for PV. Different components are dominating for all four sources. For PV, the short-term component accounts for 80% of the SD of the hourly generation time series (σ_{raw}). For wind, the mid/short-term component is dominating, for wave and load, long-term variability is most prominent and for tidal, the mid-term component accounts for 54% of σ_{raw}.

As shown in Eq. (3.7) on page 39, the variability of the net load can be explained by the weight and SD of each time series, and the correlations between them. In order to obtain a low net load variability, it is desirable with high correlations to the load and low correlation between the sources. Tidal generation is close to uncorrelated to the other sources and the load for all frequency components. PV has a strong negative correlation (between -0.83 and -0.75) to wind, wave and load for the long-term component. PV is, in contrast, positively correlated (0.42) to the load for the short-term components. Adding smaller amounts of PV will thus reduce the short-term component of the net load and wind and wave deployment will reduce the long-term component. However, after a certain point the net load variability is increased due to the relatively large SDs of output from the IRE sources. As an example, there is a short-term variability minimum when PV is supplying 2–3% of the load. For high shares of PV, the variability can be very large. Wind, wave and load are strongly correlated (0.74–0.93) to each other for the long-term components.

For each scenario (highly and fully renewable systems), both predefined mixes and optimised shares of IRE were considered. Mix 1 was a reference mix similar to today's, only slightly more PV (90% of energy from wind

[7]Note that this definition is different than that normally used. The reason for subtracting nuclear and thermal generation from the load was that we assumed the generation patterns from these sources to be fixed.

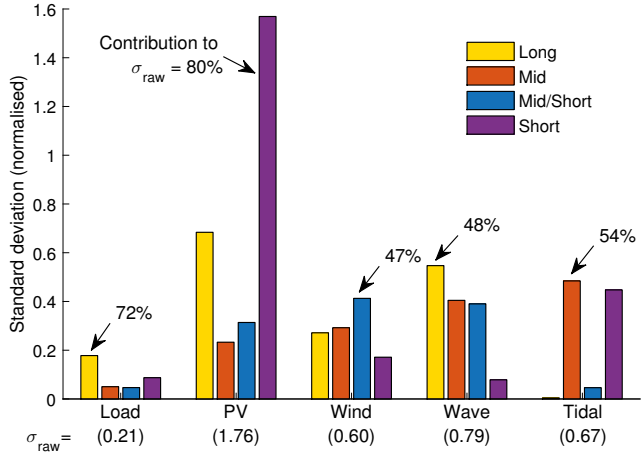

Figure 4.25. Standard deviations for the frequency components, normalised to mean load/generation. For the IRE sources, results for the 3 TWh scenarios are given (see Paper VII). The contributions to the raw σ from the dominating frequency components are indicated in the figure. The normalised σ of the raw signals are given in the horizontal axis labels.

and 10% from PV). Mix 2 was more futuristic with maximum tidal energy (3 TWh/year) and with PV, wind and wave accounting for 40%, 40% and 20% of the remaining respectively. The optimisations were carried out in order to reduce the net load SDs for each frequency component (σ_{LT}, σ_{MT}, σ_{MST} and σ_{ST}) as well as for the raw data (σ_{raw}). As can be anticipated from the SDs for and correlations between the IRE sources, the optimised mixes varied heavily depending on the objective function (Fig. 2 in Paper VII). Most mixes comprised the maximum allowed tidal generation. The mixes optimised for low σ_{raw}, σ_{LT} and σ_{ST} contained little or no PV. When σ_{MT} or σ_{MST} were minimised, PV shares of around 50% were however obtained.

The resulting net load variability is shown in Fig. 4.26. The variability is given in terms of SDs of raw data (hourly energy time series), the four frequency components and one hour step changes. As an example, let us consider the raw net load in the scenario with 20% IRE, i.e. the leftmost group of bars in Fig. 4.26. With mix 1 (90% wind and 10% PV), the SD is 7.6 GW. When the IRE shares were optimised to obtain a low σ_{raw}, the corresponding figure is 7.0 GW. Minimising σ_{MT} leads to a high share of PV, resulting in $\sigma_{raw} = 11.4$ GW. These figures should be compared to 5.9 GW for the present generation portfolio (dashed line). An increased net load variability is thus inevitable, no matter how the IRE sources are combined. The magnitude of the increase depends heavily on the IRE mix.

Two important conclusions can be drawn. Firstly, no matter how the IRE sources are combined, a large increase will be seen in σ_{MT} and σ_{MST}, i.e. fluctuations with periods ranging from two days to four months. More studies are

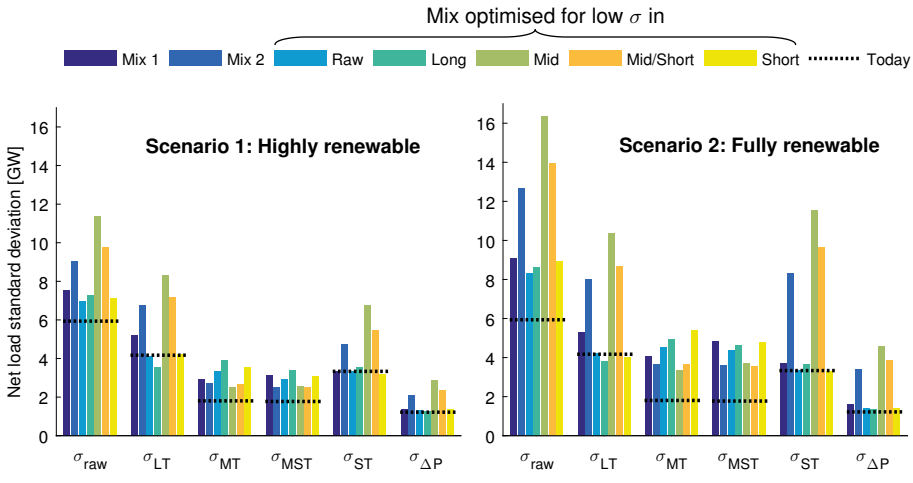

Figure 4.26. Net load standard deviations of different kind for scenario 1 and 2 compared to today's situation. The different colours indicate the mix of intermittent renewables (either predefined or optimised). Step changes are denoted ΔP.

needed in order to answer whether the remaining system, most importantly hydro power, will be able to balance these fluctuations, taking into account both technical and environmental constraints. Secondly, the improvements from using optimised mixes over mix 1 are relatively small. A small overproduction and corresponding curtailment of mix 1 can give similar net load SDs as for the optimised mixes. The system benefits from combining the sources are thus limited and consequently the IRE sources will have to compete mainly with their €/MWh costs. The "overproduction stategy" [118–120] can also be employed to obtain an even lower net load SD, e.g. similar to that today, see Fig. 4 in Paper VII.

Time series for the current situation and highly/fully renewable scenarios for year 2014 are shown in Fig. 4.27. For the latter, the IRE mixes optimised for low σ_{raw} were used. For the fully renewable scenario, the net load sometimes becomes very low and, seldomly, even negative. If export is not possible at these occasions, curtailment of IRE might be necessary. The maximum net loads are 6.5 and 9.0 GW higher than today's 44 GW for the highly and fully renewable scenarios respectively. For the reference mix (90% wind and 10% PV), the same figures are 7.0 and 11.2 GW. In other words, the maximum capacity of hydro power, import-export and any new peak capacity such as gas turbines needs to be significantly higher than the maximum utilised capacity we have seen over the last five years. The energy that needs to be provided by this additional (or currently unused) capacity is however below 1 TWh/year.

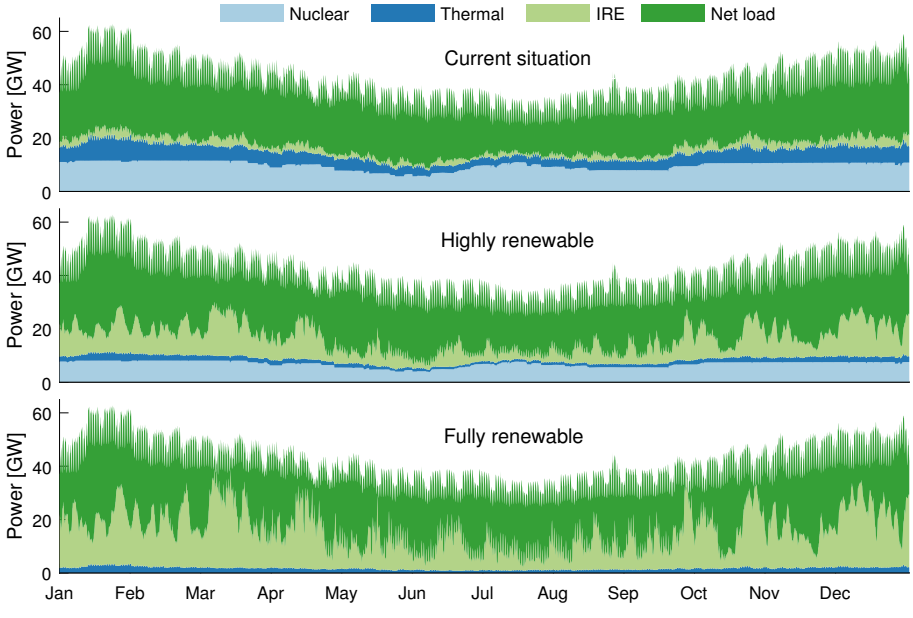

Figure 4.27. Contribution from nuclear, thermal, intermittent renewables and net load (hydro + import - export) for the year 2014. For the highly and fully renewable scenarios, results for the mixes optimised to reduce σ_{raw} are shown.

5. Concluding discussion

This thesis contains seven papers with two main foci: Papers I–IV deal with the modelling of wind power generation and forecasts and Papers V–VII with the variability of intermittent renewables.

The most important conclusion from the former group of papers is that coarse reanalyses are highly suitable for modelling, but only for generation on the power system level. Statistical post-processing is also necessary, e.g. for obtaining correct magnitudes of the high-frequency fluctuations and appropriate correlations of forecast errors. Since physical models based on reanalysis data perform so well (if used properly), linear scaling of historical generation and purely statistical models should, in my opinion, be avoided for studies of the future. If historical wind power measurements are used, a severe overestimation of the challenges and costs of wind integration may result; future WTs will have higher CFs than today's and thus give lower variability. An important reason for not using purely statistical models is that wind and load data should be synchronous.

Nevertheless, a few potential issues with the "reanalysis approach" are worth pointing out:

- Coarse reanalyses should not be used for determining mean wind speeds. The time series should rather be scaled according to mean wind speeds from highly resolved maps or estimated from energy yields of future farms.
- Coarse reanalysis data are not suitable for determining possible or likely sites for wind power deployment. Planned farms or highly resolved wind maps are better options.
- The models should always be validated with measurements. If system-level output is desired, system-level measurements should be used for tuning the model parameters.
- Statistical post-processing of the output from reanalysis based models might be necessary, e.g. if the time series are too smooth.

When developing scenarios of future farms, the most influential factors for the combined variability are the average CF and the geographical distribution. Offshore farms affect both these factors and thus require special attention. For a given set of farm characteristics, the variability is saturated relatively quickly. This implies that linear scaling of the time series are acceptable for studying e.g. different penetration levels.

In Papers I–IV, a lot of details and extra features were considered. For practical purposes, the models can be simplified with little loss of performance. In

the basic model (Paper I), it is recommended to omit time-varying air density and wind shear coefficients. In the improved model (Paper II), EOFs as predictors are, for most purposes, not worth the effort and also make the model more system specific. The search for phase angles in Papers II and III is however recommended, since this improves the representation of the varying volatility significantly.

An important conclusion from the variability studies in Papers V–VII is that intermittent renewables and the load have very different variability characteristics in different frequency bands. Wind power seasonality is highly correlated between the European countries. Fortunately, the load has a similar seasonal profile as wind power, at least for the northern countries. Even at high penetration rates, there will thus be no, or only a small, increase in the need for seasonal storage. Less fortunately, wind power in the Nordic system produces slightly less than average during the peak-load hours, so additional peak capacity (e.g. gas turbines) and/or load flexibility is likely needed for a highly renewable system.

Fluctuations with periods between a few days and a few months are dominant for wind power. Since these fluctuations are not so highly correlated between countries, interconnection can reduce the variability significantly. For the Nordic region, net load fluctuations of this type will increase strongly for high penetrations of intermittent renewables, no matter how the sources are combined. More studies, taking into account e.g. forecast uncertainties and environmental regulations, are necessary for determining whether hydro power will be able to balance these fluctuations. High-frequency variability (e.g. diurnal fluctuations and hourly step changes) are lowly correlated and not very prominent for wind power. With a suitable mix of renewables, there will be no increase in net load fluctuations on this time scale, even for a fully renewable power system.

Our work so far has been focused on modelling of renewables and studies of variability. One must therefore be careful in drawing conclusions on the impact of more wind power on the power system and possible/desirable levels of wind penetration. It can however be noted that the increase in e.g. net load step changes and required curtailments in order to not exceed certain instantaneous penetration levels are non-linear functions of the wind penetration. The challenges related to e.g. a 30% wind penetration are therefore considerably smaller than those for a 50% penetration.

According to our studies, higher CFs, more offshore wind power, more optimal geographical distribution and overproduction/curtailment would be beneficial for the system in terms of lower variability and/or lower FEs. Especially for higher wind penetrations, it might thus be rational to use policy instruments to steer the development in this direction. Before such measures are recommended, the socio-economic implications must be assessed, see suggestions for future work in the subsequent chapter.

6. Future work

Now that the model package for generating time series of wind power production and forecasts is completed, the next logical step is to focus on the implications for the power system from an increased share of wind and other intermittent renewables. The overarching goal is to answer whether a fully renewable Nordic system is feasible and to investigate suitable strategies for facilitating the integration of renewables. Two ongoing studies are however planned to first be completed: one on the possible deterioration of wind farm performance over time (similar to [23]) and one on comparing different reanalyses for modelling wind power.

A highly interesting research topic is to evaluate different strategies for mitigating the variability of renewable energy. In [39], it was shown that "advanced" WTs with lower specific rating and higher hub heights are beneficial from a system point of view and increase the revenues for wind farm owners. In this thesis, it has been demonstrated that higher CFs leads to both lower variability and lower FEs. Other possible strategies include wise combinations of renewables (see Paper VII), more offshore wind power, overproduction/curtailment and a better geographic distribution of farms, e.g. by procuring farms at strategic locations.

In order to fairly assess the abovementioned strategies from both a technical and socio-economical perspective, and to compare these to grid reinforcements, investments in new peak generation etc., a lot of work will become necessary. A transmission grid model of the studied area needs to be developed as well as models for hydropower and the future electric load. Furthermore, the electric market needs to be accounted for, perhaps through a collaboration with energy economists.

7. Summary of papers

Paper I

Modelling the Swedish wind power production using MERRA reanalysis data

Description and validation of a model for hourly, aggregated wind power generation in Sweden. The model uses MERRA reanalysis data and information on WTs and takes into account e.g. power curve smoothing and direction dependent losses. When comparing model output to measurements, the errors were low, the correlation was high and the distributions were very similar. The author collected the data, developed the model and wrote the paper.
Published in *Renewable Energy*, 76: 717–725, 2015.

Paper II

Restoring the missing high-frequency fluctuations in a wind power model based on reanalysis data

Since the above-mentioned model (Paper I) understimated the high-frequency fluctuations, a statistical supplement model was developed. By using frequency domain and machine learning techniques, the spectrum and step change distribution of the model output became very similar to measurements at the expense of a small increase in the RMS error. Furthermore, the varying volatility of the high-frequency fluctuations was captured relatively well. The author collected the data, developed the model and wrote the paper.
Published in *Renewable Energy*, 96: 784–791, 2016.

Paper III

Simulating intra-hourly wind power fluctuations on a power system level

For some power system studies, wind power data with higher resolution than hourly are desirable. This paper describes and validates a model for simulation of intra-hourly fluctuations, given that hourly data are available. It was shown that a model trained on data from one power system can be used for simulating realistic fluctuations for another system, e.g. in terms of distribution and varying volatility of the deviations from hourly means. The author collected the data, developed the model in collaboration with the co-authors and wrote the paper.
Resubmitted after revisions to *Wind Energy*, August 2016.

Paper IV

A new approach to obtain synthetic wind power forecasts for integration studies

For higher penetration levels, the forecast errors of wind power can be significant and must be accounted for in wind integration studies. This paper describes a new approach for obtaining "synthetic" wind power forecasts with a horizon of up to one week or more. Unlike earlier methodologies, freely available, global *reforecasts* (from the GEFS dataset) were used. In order to produce realistic forecasts, some statistical post-processing was employed. Validation with aggregated observations from Belgium and three individual farms in Sweden proved the adequacy of the approach. The author collected some of the data, developed the model and wrote the paper.

Resubmitted after revisions to *Energies* September, 2016.

Paper V

Variability Assessment and Forecasting of Renewables: A Review for Solar, Wind, Wave and Tidal Resources

Very large amounts of research has been conducted on the variability and forecastability of renewable energy sources, in particular for solar and wind power. The questions studied and methods employed differs between the sources. This review attempts to summarize and compare the research done so far, and points out research gaps and areas of possible learning between the fields. The author collected and read the references related to wind energy and wrote the corresponding text. The author also contributed to merging the text and in discussing the scope and results.

Published in *Renewable & Sustainable Energy Reviews*, 44: 356-375, 2015.

Paper VI

Correlation between wind power generation in the European countries

The correlations of wind power generation in different countries are important since they give information on the benefits of interconnection. Country-wise time series were generated for a European scenario for year 2020. It was shown that the correlation decreases in an exponential fashion with distance. In order to deepen the analysis, the time series were band-pass filtered. Long-term components (seasonal variations) were highly correlated for most countries, but fluctuations with shorter periods were less correlated. The author collected the data, performed the analyses and wrote the paper.

Published in *Energy*, 114: 663–670, 2016.

Paper VII

Net load variability in the Nordic countries with a highly or fully renewable power system

Scenarios for the Nordic power system were developed in which intermittent renewables (PV, wind, wave and tidal) replaced fossil and nuclear generation. Different mixes, predefined or optimised, of the renewables were considered. The seasonal and diurnal fluctuations that hydropower needs to balance do not necessarily increase, even for a fully renewable system. Fluctuations with intermediate periods (days to months) will however inevitable increase significantly, no matter how the renewables are combined. The author modelled wind power generation, did the analyses of combinations of sources and wrote most of the paper.

Resubmitted after revisions to *Nature Energy*, August 2016.

8. Svensk sammanfattning

Under det senaste årtiondet har stora mängder vindkraft installerats i många länder. Denna utveckling har drivits av t.ex. lägre kostnader och miljöhänsyn och förväntas fortsätta i framtiden. När vindkraft och andra intermittenta förnybara källor (IRE) börjar försörja en större andel av elförbrukningen ifrågasätts ibland om den inneboende variabiliteten och oförutsägbarheten hos dessa kan hanteras. För att studera påverkan på kraftsystemet från en högre andel IRE utförs integrationsstudier. Realistiska tidsserier av vindkraftsproduktion och -prognoser är viktiga beståndsdelar i dessa studier. Tidsserierna ska helst vara synkrona med förbrukningsdata, vara flera år långa och ha en tillräckligt hög tidsupplösning.

Modellpaketet som beskrivs och utvärderas i uppsatserna I–IV ger en heltäckande metodik för att simulera vinddata för kraftsystemstudier. I uppsats I beskrivs den grundläggande modellen för att generera tidsserier av vindkraftproduktion utifrån information om vindkraftsparker samt grovupplösta meteorologiska data. Genom att jämföra den modellerade, timvisa produktionen för Sverige med uppmätta data visar vi att metoden fungerar tillfredställande; korrelationen är 0.98 och fördelningarna mycket likartade. Storleken på högfrekventa fluktuationer underskattas dock vilket motiverade oss att utveckla en statistisk efterbehandlingsmodul vilken beskrivs och utvärderas i uppsats II. För vissa typer av studier är det önskvärt med data av högre upplösning än en timme. Eftersom inga sådana mätningar finns tillgängliga på aggregerad nivå i Sverige samt att fritt tillgängliga meteorologiska modeller som bäst har upplösningen en timme utvecklades en statistisk modell för att kunna simulera stokastiska inomtimmen-variationer. Vi visar i uppsats III att modellen ger bra resultat även om den tränas på data från ett kraftsystem och appliceras på ett annat. Slutligen presenteras i uppsats IV en ny metodik för att generera "syntetiska" vindkraftsprognoser baserade på fritt tillgängliga data från grovupplösta meteorologiska modeller.

Den viktigaste slutsatsen från ovanstående studier är att grovupplösta meteorologiska återanalyser är högst lämpliga, men bara för att modellera den aggregerade produktionen i ett tillräckligt stort område, exempelvis ett land. En viss statistisk efterbehandling är också nödvändigt för att t.ex. få en bra återgivning av högfrekventa fluktuationer och lämpliga korrelationer av prognosfel. Fördelar med vår metodik inkluderar i) en fysisk koppling till väderdata och vindkraftverkens egenskaper, ii) över 30 år tidsserier med upplösning 5–60 minuter, iii) fritt och globalt tillgängliga indata samt iv) beräkningstider i storleksordningen minuter. I denna avhandling argumenterar jag för att vår

metodik generellt är att föredra framför att använda t.ex. rent statistiska modeller eller linjär skalning av historiska mätningar. Ett antal fallgropar finns dock vid användningen av återanalyser.

I variabilitetsstudierna som presenteras i uppsatserna V–VII studeras elproduktion från sol, vind, havsvågor och tidvatten. En översikt av forskningsläget (uppsats V) visar, föga förvånande, att forskningen kommit längst inom sol- och vindkraft. En mängd studier finns av variabiliteten och prognostiserbarheten för dessa källor, både separat och i kombination. Fördelar finns av att kombinera IRE och sprida produktionen geografiskt eftersom variabiliteten då minskar, men mcket få studier analyserar alla fyra källorna tillsammans. Det skulle kunna vara fördelaktigt för de olika forskningsfälten att lära av varandra samt att använda mer konsistenta metoder och mått.

En viktig slutsats från uppsatserna VI och VII är att de olika IRE-källorna samt elförbrukningen har väldigt olika karaktäristik i olika frekvensband. Beroende på storleken på och korrelationen av dessa fluktuationer kommer olika tidsskalor bli mer eller mindre utmanande att balansera i ett kraftsystem med mycket förnybart. Vindkraften har liknande säsongsmönster i större delen av Europa vilket innebär att även en helt sammankopplad kontinent skulle ha betydande variationer av produktionen på denna tidsskala. Korrelationerna minskar dock ju kortare tidsskalor man studerar. Fluktuationerna med perioder kortare än två dygn är t.ex. mycket lågt korrelerade även för närliggande länder och variabiliteten minskar följdaktligen effektivt genom sammankoppling. I ett helt förnybart nordiskt kraftsystem kommer det bli ingen eller endast en liten ökning av behovet för säsongslagring givet en lämplig mix av IRE. Fluktuationer med perioder mellan några dagar och några månader är dominanta för vindkraften och variationerna av denna typ som vattenkraften behöver balansera kommer att öka betydligt, oavsett hur de förnybara källorna kombineras. Med en lämplig mix av förnybart kommer däremot variationerna inom dygnet inte öka överhuvudtaget ens för ett helt förnybart system. En lovande strategi för att minska variabiliteten är att bygga mer IRE och sedan strypa produktionen vid tillfällen när den inte behövs.

Enligt våra studier leder högre kapacitetsfaktorer och mer havsbaserad vindkraft till lägre variabilitet och mindre prognosfel. I synnerhet för höga andelar vindkraft skulle det därför kunna vara rationellt att använda styrmedel för att få till stånd en sådan utveckling. En intressant inriktning för framtida forskning är att jämföra olika strategier för att minska variabiliteten hos de intermittenta källorna. För att kunna besvara om dessa är samhällsekonomiskt försvarbara och hur väl de står sig i konkurrens med t.ex. ökad utbyggnad av balanskraft och utlandsförbindelser krävs samarbeten med vattenkrafts- och elmarknadsforskare.

References

[1] P. W. Carlin, A. S. Laxson, and E. B. Muljadi, "The history and state of the art of variable-speed wind turbine technology," *Wind Energy*, vol. 6, no. 2, 2003.

[2] J. Olauson, "Wind power and natural disasters," Licentiate Thesis, Uppsala University, 2014, ISSN 0349-8352; 337-14L.

[3] "Global wind energy outlook 2014," Global Wind Energy Council, Tech. Rep., 2014.

[4] I. Nohlgren, S. Herstad Sväd, M. Jansson, and J. Rodin, "El från nya och framtida anläggningar 2014," Elforsk, Tech. Rep. 14:40, 2014.

[5] H. Holttinen, P. Meibom, A. Orths, B. Lange, M. O'Malley, J. O. Tande, A. Estanqueiro, E. Gomez, L. Söder, G. Strbac, J. C. Smith, and F. van Hulle, "Impacts of large amounts of wind power on design and operation of power systems, results of IEA collaboration," *Wind Energy*, vol. 14, no. 2, pp. 179–192, 2011.

[6] J. Kiviluoma, H. Holttinen, D. Weir, R. Scharff, L. Söder, N. Menemenlis, N. A. Cutululis, I. Danti Lopez, E. Lannoye, A. Estanqueiro, E. Gomez-Lazaro, Q. Zhang, J. Bai, Y.-H. Wan, and M. Milligan, "Variability in large-scale wind power generation," *Wind Energy*, In print, DOI: 10.1002/we.1942.

[7] L. Hirth, "The market value of variable renewables: The effect of solar wind power variability on their relative price," *Energy economics*, vol. 38, pp. 218–236, Jul. 2013.

[8] P. C. Bhagwat, L. J. de Vries, and B. F. Hobbs, "Expert survey on capacity markets in the US: Lessons for the EU," *Utilities Policy*, vol. 38, pp. 11–17, Feb. 2016.

[9] M. Huber, D. Dimkova, and T. Hamacher, "Integration of wind and solar power in Europe: Assessment of flexibility requirements," *Energy*, vol. 69, pp. 236–246, May 2014.

[10] K. Schaber, F. Steinke, and T. Hamacher, "Transmission grid extensions for the integration of variable renewable energies in Europe: Who benefits where?" *Energy Policy*, vol. 43, pp. 123–135, 2012.

[11] P. Tielens and D. Van Hertem, "The relevance of inertia in power systems," *Renewable and Sustainable Energy Reviews*, vol. 55, pp. 999–1009, Mar. 2016.

[12] A. K. Pathak, M. P. Sharma, and M. Bundele, "A critical review of voltage and reactive power management of wind farms," *Renewable and Sustainable Energy Reviews*, vol. 51, pp. 460–471, Nov. 2015.

[13] J. Fortmann, R. Pfeiffer, E. Haesen, F. van Hulle, F. Martin, H. Urdal, and S. Wachtel, "Fault-ride-through requirements for wind power plants in the ENTSO-E network code on requirements for generators," *IET Renewable Power Generation*, vol. 9, no. 1, pp. 18–24, Jan. 2015.

[14] L. Söder, "På väg mot en elförsörjning baserad på enbart förnybar el i Sverige - En studie om behov av reglerkraft och överföringskapacitet (Version 4.0)," KTH, Tech. Rep., 2014.
[15] J. Bruce, S. Holmer, A. Badano, L. Söder, S. Larsson, N. Dahlbäck, J. Bladh, J. Lönnberg, L. Göransson, B. Rydén, H. Sköldberg, T. Unger, and S. Montin, "Reglering av kraftsystemet med ett stort inslag variabel produktion," North European Power Perspectives (NEPP), Tech. Rep., 2016.
[16] IRENA, "The power to change: Solar and wind cost reduction potential to 2025," International Renewable Energy Agency, Tech. Rep. ISBN 978-92-95111-97-4 (PDF), 2016.
[17] C. Potter, D. Lew, J. McCaa, S. Cheng, S. Eichelberger, and E. Grimit, "Creating the dataset for the western wind and solar integration study (U.S.A.)," *Wind Engineering*, vol. 32, pp. 325–338, Jun. 2008.
[18] H. Holttinen, "Expert group report on recommended practices: 16. Wind integration studies, 1st edition," IEA Wind, Tech. Rep., Sep. 2013.
[19] H. Holttinen, M. O'Malley, J. Dillon, D. Flynn, A. Keane, H. Abildgaard, and L. Söder, "Steps for a complete wind integration study," in *46th Hawaii International Conference on System Sciences*, 2013.
[20] H. Holttinen, J. Kiviluoma, A. Robitaille, N. Cutululis, A. Orths, F. van Hulle, I. Pineda, B. Lange, M. O'Malley, J. Dillon, E. Carlini, C. Vergine, J. Kondoh, M. Gibescu, J. Tande, A. Estanqueiro, E. Gomez, L. Söder, J. Smith, M. Milligan, and D. Lew, "Design and operation of power systems with large amounts of wind power, Final summary report, IEA WIND Task 25, Phase three 2012–2014," VTT Technology 75, Espoo, Finland, Tech. Rep., 2013.
[21] M. M. Rienecker, M. J. Suarez, R. Gelaro, R. Todling, J. Bacmeister, E. Liu, M. G. Bosilovich, S. D. Schubert, L. Takacs, G.-K. Kim, S. Bloom, J. Chen, D. Collins, A. Conaty, A. da Silva, W. Gu, J. Joiner, R. D. Koster, R. Lucchesi, A. Molod, T. Owens, S. Pawson, P. Pegion, C. R. Redder, R. Reichle, F. R. Robertson, A. G. Ruddick, M. Sienkiewicz, and J. Woollen, "MERRA: NASA's modern-era retrospective analysis for research and applications," *Journal of Climate*, vol. 24, pp. 3624–3648, Jul. 2011.
[22] D. P. Dee, S. M. Uppala, A. J. Simmons, P. Berrisford, P. Poli, S. Kobayashi, U. Andrae, M. A. Balmaseda, G. Balsamo, P. Bauer, P. Bechtold, A. C. M. Beljaars, L. van de Berg, J. Bidlot, N. Bormann, C. Delsol, R. Dragani, M. Fuentes, A. J. Geer, L. Haimberger, S. B. Healy, H. Hersbach, E. V. Hólm, L. Isaksen, P. Kållberg, M. Köhler, M. Matricardi, A. P. McNally, B. M. Monge-Sanz, J.-J. Morcrette, B.-K. Park, C. Peubey, P. de Rosnay, C. Tavolato, J.-N. Thépaut, and F. Vitart, "The ERA-Interim reanalysis: configuration and performance of the data assimilation system," *Quarterly Journal of the Royal Meteorological Society*, vol. 137, no. 656, 2011.
[23] I. Staffell and R. Green, "How does wind farm performance decline with age?" *Renewable Energy*, vol. 66, pp. 775–786, Jun. 2014.
[24] M. L. Kubik, D. J. Brayshaw, P. J. Coker, and J. F. Barlow, "Exploring the role of reanalysis data in simulating regional wind generation variability over northern ireland," *Renewable energy*, vol. 57, pp. 558–561, Sep. 2013.
[25] D. Heide, L. von Bremen, M. Greiner, C. Hoffmann, M. Speckmann, and S. Bofinger, "Seasonal optimal mix of wind and solar power in a future, highly

renewable Europe," *Renewable Energy*, vol. 35, no. 11, pp. 2483–2489, Nov. 2010.
[26] G. B. Andresen, A. A. Søndergaard, and M. Greiner, "Validation of Danish wind time series from a new global renewable energy atlas for energy system analysis," *Energy*, vol. 93, Part 1, pp. 1074–1088, Dec. 2015.
[27] D. R. Drew, D. J. Cannon, D. J. Brayshaw, J. F. Barlow, and P. J. Coker, "The impact of future offshore wind farms on wind power generation in Great Britain," *Resources*, vol. 4, no. 1, pp. 155–171, Mar. 2015.
[28] I. Staffell and S. Pfenninger, "Using bias-corrected reanalysis to simulate current and future wind power output," *Energy*, vol. 114, pp. 1224–1239, Nov. 2016.
[29] F. van Hulle, "Integrating wind: Developing Europe's power market for the large-scale integration of wind power," TradeWind, Tech. Rep., 2009.
[30] S. Rose and J. Apt, "Generating wind time series as a hybrid of measured and simulated data," *Wind Energy*, vol. 15, pp. 699–715, 2012.
[31] L. Söder, "Integration study of small amounts of wind power in the power system," KTH, Tech. Rep. Trita-EES-9401, 1994.
[32] J. Apt, "The spectrum of power from wind turbines," *Journal of Power Sources*, vol. 169, no. 2, pp. 369–374, Jun. 2007.
[33] W. C. Skamarock, "Evaluating mesoscale NWP models using kinetic energy spectra," *Monthly Weather Review*, vol. 132, pp. 3019–3032, 2004.
[34] R. J. Davy, M. J. Woods, C. J. Russell, and P. A. Coppin, "Statistical downscaling of wind variability from meteorological fields," *Boundary-Layer Meteorology*, vol. 135, no. 1, pp. 161–175, Jan. 2010.
[35] N. Ellis, R. Davy, and A. Troccoli, "Predicting wind power variability events using different statistical methods driven by regional atmospheric model output," *Wind Energy*, vol. 18, no. 9, pp. 1611–1628, Sep. 2015.
[36] C. M. St. Martin, J. K. Lundquist, and M. A. Handschy, "Variability of interconnected wind plants: correlation length and its dependence on variability time scale," *Environmental Research Letters*, vol. 10, no. 4, p. 044004, 2015.
[37] H. Louie, "Correlation and statistical characteristics of aggregate wind power in large transcontinental systems," *Wind Energy*, vol. 17, no. 6, pp. 793–810, 2014.
[38] E. Fertig, J. Apt, P. Jaramillo, and W. Katzenstein, "The effect of long-distance interconnection on wind power variability," *Environmental research letters*, vol. 7, pp. 1–6, 2012.
[39] L. Hirth and S. Müller, "System-friendly wind power: How advanced wind turbine design can increase the economic value of electricity generated through wind power," *Energy Economics*, vol. 56, pp. 51–63, May 2016.
[40] P. Sørensen, N. A. Cutululis, A. Vigueras-Rodríguez, H. Madsen, P. Pinson, L. E. Jensen, J. Hjerrild, and M. Donovan, "Modelling of power fluctuations from large offshore wind farms," *Wind Energy*, vol. 11, pp. 29–43, 2008.
[41] H. Holttinen, "Hourly wind power variations in the Nordic countries," *Wind Energy*, vol. 8, no. 2, pp. 173–195, 2005.
[42] J. Olsson, L. Skoglund, F. Carlsson, and L. Bertling, "Future wind power production variations in the Swedish power system," in *Innovative Smart Grid*

Technologies Conference Europe (ISGT Europe), 2010 IEEE PES, 2010, pp. 1–7.

[43] R. Billinton, H. Chen, and R. Ghajar, "Time-series models for reliability evaluation of power systems including wind energy," *Microelectronics Reliability*, vol. 36, no. 9, pp. 1253–1261, Sep. 1996.

[44] G. Papaefthymiou and B. Klöckl, "MCMC for wind power simulation," *IEEE Transactions on Energy Conversion*, vol. 23, pp. 234–240, 2008.

[45] M. Milligan, E. Ela, D. Lew, D. Corbus, Y.-h. Wan, and B. Hodge, "Assessment of simulated wind data requirements for wind integration studies," *IEEE Transactions on Sustainable Energy*, vol. 3, pp. 620–626, Oct. 2012.

[46] S. Eriksson, "Direct driven generators for vertical axis wind turbines," PhD Thesis, Uppsala University, 2008, ISBN 978-91-554-7264-1.

[47] F. Bülow, "A generator perspective on vertical axis wind turbines," PhD Thesis, Uppsala University, 2013, ISBN 978-91-554-8642-6.

[48] A. Goude, "Fluid mechanics of vertical axis turbines: Simulations and model development," PhD Thesis, Uppsala University, 2012, ISBN 978-91-554-8539-9.

[49] E. Möllerström, F. Ottermo, J. Hylander, and H. Bernhoff, "Noise emission of a 200 kW vertical axis wind turbine," *Energies*, vol. 9, no. 1, 2016.

[50] J. Kjellin, "Vertical axis wind turbines: Electrical system and experimental results," PhD Thesis, Uppsala University, 2012, ISBN 978-91-554-8496-5.

[51] S. Eriksson, J. Kjellin, and H. Bernhoff, "Tip speed ratio control of a 200 kW VAWT with synchronous generator and variable DC voltage," *Energy Science & Engineering*, vol. 1, no. 3, 2013.

[52] J. Kjellin, F. Bülow, S. Eriksson, P. Deglaire, M. Leijon, and H. Bernhoff, "Power coefficient measurement on a 12 kW straight bladed vertical axis wind turbine," *Renewable Energy*, vol. 36, no. 11, pp. 3050–3053, 2011.

[53] S. Eriksson, H. Bernhoff, and M. Bergkvist, "Design of a unique direct driven PM generator adapted for a telecom tower wind turbine," *Renewable Energy*, vol. 44, pp. 453–456, 2012.

[54] J. Arnqvist, "Mean wind and turbulence conditions in the boundary layer above forests," PhD Thesis, Uppsala University, 2015, ISBN 978-91-554-9123-9.

[55] P. Thorsson, S. Söderberg, and H. Bergström, "Modelling atmospheric icing: A comparison between icing calculated with measured meteorological data and NWP data," *Cold Regions Science and Technology*, vol. 119, pp. 124–131, Nov. 2015.

[56] K. Nilsson, S. Ivanell, K. S. Hansen, R. Mikkelsen, J. N. Sørensen, S.-P. Breton, and D. Henningson, "Large-eddy simulations of the Lillgrund wind farm," *Wind Energy*, vol. 18, no. 3, pp. 449–467, Mar. 2015.

[57] C. B. Field, V. R. Barros, D. J. Dokken, K. J. Mach, M. D. Mastrandrea, T. E. Bilir, M. Chatterjee, K. L. Ebi, Y. O. Estrada, R. C. Genova, B. Girma, E. S. Kissel, A. N. Levy, S. MacCracken, P. R. Mastrandrea, and L. L. White, "Summary for policymakers," in *Climate Change 2014: Impacts, Adaptation, and Vulnerability. Part A: Global and Sectoral Aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge, United Kingdom and New York, NY,

USA: Cambridge University Press, 2014, pp. 1–32.
[58] "Svenska kraftnät, Statistik," Available online: http://www.svk.se/aktorsportalen/elmarknad/statistik/ (Accessed: 2016-04-13).
[59] "Svensk Vindenergi, Statistik om vindkraft," Available online: http://www.vindkraftsbranschen.se/statistik/ (Accessed: 2016-03-16).
[60] "Energimyndigheten, Marknadsstatistik," Available online: http://www.energimyndigheten.se/fornybart/elcertifikatsystemet/marknadsstatistik/ (Accessed: 2016-03-16).
[61] A. Molod, L. Takacs, M. Suarez, and J. Bacmeister, "Development of the GEOS-5 atmospheric general circulation model: evolution from merra to merra2," *Geoscientific Model Development*, vol. 8, no. 5, pp. 1339–1356, 2015.
[62] S. Kobayashi, Y. Ota, Y. Harada, A. Ebita, M. Moriya, H. Onoda, K. Onogi, H. Kamahori, C. Kobayashi, H. Endo, K. Miyaoka, and K. Takahashi, "The JRA-55 Reanalysis: General Specifications and Basic Characteristics," *Journal of the Meteorological Society of Japan. Ser. II*, vol. 93, no. 1, pp. 5–48, 2015.
[63] S. Liléo, E. Berge, O. Undheim, R. Klinkert, and R. E. Bredesen, "Long-term correction of wind measurements - state-of-the-art, guidelines and future work," Elforsk report 13:18, Tech. Rep., Jan. 2013.
[64] E. Sharp, P. Dodds, M. Barrett, and C. Spataru, "Evaluating the accuracy of CFSR reanalysis hourly wind speed forecasts for the UK, using in situ measurements and geographical information," *Renewable Energy*, vol. 77, pp. 527–538, May 2015.
[65] D. Carvalho, A. Rocha, M. Gómez-Gesteira, and C. Silva Santos, "Comparison of reanalyzed, analyzed, satellite-retrieved and NWP modelled winds with buoy data along the Iberian Peninsula coast," *Remote Sensing of Environment*, vol. 152, pp. 480–492, Sep. 2014.
[66] S. Saha, S. Moorthi, H.-L. Pan, X. Wu, J. Wang, S. Nadiga, P. Tripp, R. Kistler, J. Woollen, D. Behringer, H. Liu, D. Stokes, R. Grumbine, G. Gayno, J. Wang, Y.-T. Hou, H.-Y. Chuang, H.-M. H. Juang, J. Sela, M. Iredell, R. Treadon, D. Kleist, P. Van Delst, D. Keyser, J. Derber, M. Ek, J. Meng, H. Wei, R. Yang, S. Lord, H. Van Den Dool, A. Kumar, W. Wang, C. Long, M. Chelliah, Y. Xue, B. Huang, J.-K. Schemm, W. Ebisuzaki, R. Lin, P. Xie, M. Chen, S. Zhou, W. Higgins, C.-Z. Zou, Q. Liu, Y. Chen, Y. Han, L. Cucurull, R. W. Reynolds, G. Rutledge, and M. Goldberg, "The NCEP climate forecast system reanalysis," *Bulletin of the American Meteorological Society*, vol. 91, no. 8, pp. 1015–1057, Apr. 2010.
[67] L. Reichenberg, F. Johnsson, and M. Odenberger, "Dampening variations in wind power generation - the effect of optimizing geographic location of generating sites," *Wind Energy*, vol. 17, no. 11, pp. 1631–1643, Nov. 2014.
[68] D. J. Cannon, D. J. Brayshaw, J. Methven, P. J. Coker, and D. Lenaghan, "Using reanalysis data to quantify extreme wind power generation statistics: A 33 year case study in Great Britain," *Renewable Energy*, vol. 75, pp. 767–778, Mar. 2015.
[69] J. Carta, S. Velázquez, and P. Cabrera, "A review of measure-correlate-predict (MCP) methods used to estimate long-term wind characteristics at a target site," *Renewable and Sustainable Energy Reviews*, vol. 27, pp. 362–400, Nov. 2013.

[70] H. Bergström, "Boundary-layer modelling for wind climate estimates," *Wind engineering*, vol. 25, no. 5, pp. 289–299, 2001.
[71] L. Enger, "Simulation of dispersion in a moderately complex terrain. Part A. The fluid dynamic model," *Atmospheric Environment*, vol. 24A, pp. 2431–2446, 1990.
[72] A. Andrén, "Evaluation of a turbulence closure scheme suitable for air pollution applications," *Journal of Applied Meteorology*, vol. 29, pp. 224–239, 1990.
[73] W. Skamarock, J. Klemp, J. Dudhia, D. Gill, M. Barker, K. Duda, Y. Huang, W. Wang, and J. Powers, "A description of the advanced research WRF version 3," National Center for Atmospheric Research, Tech. Rep., 2008.
[74] R. Swinbank, M. Kyouda, P. Buchanan, L. Froude, T. M. Hamill, T. D. Hewson, J. H. Keller, M. Matsueda, J. Methven, F. Pappenberger, M. Scheuerer, H. A. Titley, L. Wilson, and M. Yamaguchi, "The TIGGE project and its achievements," *Bulletin of the American Meteorological Society*, vol. 97, pp. 49–67, 2016.
[75] T. M. Hamill, G. T. Bates, J. S. Whitaker, D. R. Murray, M. Fiorino, T. J. Galarneau, Y. Zhu, and W. Lapenta, "NOAA's second-generation global medium-range ensemble reforecast dataset," *Bulletin of the American Meteorological Society*, vol. 94, no. 10, pp. 1553–1565, Feb. 2013.
[76] S. Emeis, *Wind Energy Meteorology*. Berlin, Germany: Springer, 2013.
[77] A. Betz, *Windenergie und ihre ausnutzung durch windmüllen*. Göttingen, Germany: Vandenhoeck and Ruprecht, 1926.
[78] "Power performance measurements of electricity producing wind turbines," International Electrotechnical Commission, Geneva, Switzerland, Tech. Rep. IEC 61400-12-1:2005(E), 2005.
[79] M. L. Thøgersen, M. Motta, T. Sørensen, and P. Nielsen, "Measure-correlate-predict methods: Case studies and software implementation," in *EWEA Conference Proceedings*. EMD International A/S, 2007.
[80] Y.-H. Wan, "Analysis of wind power ramping behavior in ERCOT," National Renewable Energy Laboratory, Tech. Rep. NREL/TP-5500-49218, 2011.
[81] E. D. Stoutenburg, N. Jenkins, and M. Z. Jacobson, "Power output variations of co-located offshore wind turbines and wave energy converters in California," *Renewable Energy*, vol. 35, no. 12, pp. 2781–2791, Dec. 2010.
[82] P. Sorensen, N. Cutululis, A. Vigueras-Rodriguez, L. Jensen, J. Hjerrild, M. Donovan, and H. Madsen, "Power fluctuations from large wind farms," *IEEE Transactions on Power Systems*, vol. 22, no. 3, pp. 958–965, 2007.
[83] P. Nørgaard and H. Holttinen, "A multi-turbine power curve approach," in *Proceedings of Nordic Wind Power Conference*, Gothenburg, Sweden, Mar. 2004.
[84] R. Scharff, "Design of electricity markets for efficient balancing of wind power generation," PhD Thesis, KTH, 2015, ISBN 978-91-7595-652-7.
[85] U. Focken, M. Lange, K. Mönnich, H.-P. Waldl, H. G. Beyer, and A. Luig, "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*, vol. 90,

no. 3, pp. 231–246, Mar. 2002.
[86] C. Wan, Z. Xu, P. Pinson, Z. Y. Dong, and K. P. Wong, "Probabilistic forecasting of wind power generation using extreme learning machine," *IEEE Transactions on Power Systems*, vol. 29, no. 3, pp. 1033–1044, May 2014.
[87] A. M. Foley, P. G. Leahy, A. Marvuglia, and E. J. McKeogh, "Current methods and advances in forecasting of wind power generation," *Renewable Energy*, vol. 37, no. 1, pp. 1–8, 2012.
[88] G. Kariniotakis, I. Marti, D. Casas, P. Pinson, T. S. Nielsen, H. Madsen, G. Giebel, J. Usaola, I. Sanchez, A. M. Palomares, R. Brownsword, J. Tambke, U. Focken, M. Lange, P. Pouka, G. Kallos, C. Lac, G. Sideratos, and G. Descombes, "What performance can be expected by short-term wind power prediction models depending on site characteristics?" in *In CD-Rom Proceedings, European Wind Energy Conference EWEC*. European Wind Energy Association (EWEA), 2005.
[89] A. Boone, "Simulation of short-term wind speed forecast errors using a multi-variate ARMA(1,1) time-series model," Master thesis, KTH, Stockholm, Sweden, 2005.
[90] B.-M. Hodge, H. Holttinen, S. Sillanpää, E. Gómez-Lázaro, R. Scharff, L. Söder, X. Larsén, G. Giebel, D. Flynn, J. Dobschinski, D. Lew, and M. Milligan, "Wind power forecasting error distributions: An international comparison," in *11th Annual International Workshop on Large-Scale Integration of Wind Power into Power Systems as well as on Transmission Networks for Offshore Wind Power Plants Conference*, Lisbon, Portugal, Nov. 2012.
[91] K. J. Beven, *Environmental modelling: an uncertain future?* Abingdon, United Kingdom: Routledge, 2009.
[92] J. Nelder and R. Mead, "A simplex method for function minimization," *Computer Journal*, vol. 7, pp. 308–313, 1965.
[93] J. Lagarias, J. Reeds, M. Wright, and P. Wright, "Convergence properties of the Nelder–Mead simplex method in low dimensions," *SIAM Journal on Optimization*, vol. 9, no. 1, pp. 112–147, Jan. 1998.
[94] K. McKinnon, "Convergence of the Nelder–Mead simplex method to a nonstationary point," *SIAM Journal on Optimization*, vol. 9, no. 1, pp. 148–158, Jan. 1998.
[95] L. Söder, "Simulation of wind speed forecast errors for operation planning of multiarea power systems," in *2004 International Conference on Probabilistic Methods Applied to Power Systems*, Sep. 2004, pp. 723–728.
[96] R. W. Preisendorfer and C. D. Mobley, *Principal component analysis in meteorology and oceanography*. Amsterdam, the Netherlands: Elsevier, 1988.
[97] A. Navarra and V. Simoncini, *A Guide to Empirical Orthogonal Functions for Climate Data Analysis*. Dordrecht, the Netherlands: Springer, 2010.
[98] P. D. Welch, "The use of fast fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms," *IEEE Trans. Audio and Electroacoust.*, vol. AU-15, pp. 70–73, June 1967.
[99] S. W. Smith, *The scientist and engineer's guide to digital signal processing*, 2nd ed. San Diego, CA: California Technical Publishing, 1999, Available

online: http://www.DSPguide.com.
[100] W. Katzenstein, E. Fertig, and J. Apt, "The variability of interconnected wind plants," *Energy Policy*, vol. 38, pp. 4400–4410, Aug. 2010.
[101] D. Lee and R. Baldick, "Future wind power scenario synthesis through power spectral density analysis," *IEEE Transactions on Smart Grid*, vol. 5, pp. 490–500, Jan. 2014.
[102] Y. Li, K. Xie, and B. Hu, "Copula-ARMA model for multivariate wind speed and its applications in reliability assessment of generating systems," *Journal of Electrical Engineering and Technology*, vol. 8, no. 3, pp. 421–427, May 2013.
[103] A. Sturt and G. Strbac, "Time series modelling of power output for large-scale wind fleets," *Wind Energy*, vol. 14, no. 8, pp. 953–966, 2011.
[104] P. E. de Mello, N. Lu, and Y. Makarov, "An optimized autoregressive forecast error generator for wind and load uncertainty study," *Wind Energy*, vol. 14, no. 8, pp. 967–976, Nov. 2011.
[105] R. A. Berk, *Statistical learning from a regression perspective*, 1st ed., ser. Springer Series in Statistics. New York, N.Y: Springer, 2008.
[106] L. Breiman, J. Friedman, C. J. Stone, and R. A. Olshen, *Classification and regression trees*. Monterey, CA: Wadsworth press, Jan. 1984.
[107] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, Oct. 2001.
[108] V. Bulaevskaya, S. Wharton, A. Clifton, G. Qualley, and W. O. Miller, "Wind power curve modeling in complex terrain using statistical models," *Journal of Renewable and Sustainable Energy*, vol. 7, no. 1, p. 013103, Jan. 2015.
[109] J. H. Friedman, "Greedy function approximation: A gradient boosting machine," *The Annals of Statistics*, vol. 29, no. 5, pp. 1189–1232, 2001.
[110] A. Kusiak and Z. Zhang, "Short-horizon prediction of windpower: A data-driven approach," *IEEE Transactions on Energy Conversion*, vol. 25, no. 4, pp. 1112–1122, Dec. 2010.
[111] M. Olsson, M. Perninge, and L. Söder, "Modeling real-time balancing power demands in wind power systems using stochastic differential equations," *Electric Power Systems Research*, vol. 80, pp. 966–974, Aug. 2010.
[112] J. P. Deane, G. Drayton, and B. P. Ó Gallachóir, "The impact of sub-hourly modelling in power systems with significant levels of renewable generation," *Applied Energy*, vol. 113, pp. 152–158, Jan. 2014.
[113] N. Troy, D. Flynn, and M. O'Malley, "The importance of sub-hourly modeling with a high penetration of wind generation," in *2012 IEEE Power and Energy Society General Meeting*, Jul. 2012, pp. 1–6.
[114] A. Keane, M. Milligan, C. J. Dent, B. Hasche, C. D'Annunzio, K. Dragoon, H. Holttinen, N. Samaan, L. Soder, and M. O'Malley, "Capacity value of wind power," *IEEE Transactions on Power Systems*, vol. 26, no. 2, pp. 564–572, May 2011.
[115] L. Hirth, "The market value of wind energy - thermal versus hydro power systems," Energiforsk, Tech. Rep. 2016:276, 2016.
[116] "Nordpool, Historical market data," Available online: http://www.nordpoolspot.com/historical-market-data/ (Accessed: 2016-07-16).
[117] E. V. Mc Garrigle, J. P. Deane, and P. G. Leahy, "How much wind energy will

be curtailed on the 2020 Irish power system?" *Renewable Energy*, vol. 55, pp. 544–553, Jul. 2013.
[118] D. Heide, M. Greiner, L. von Bremen, and C. Hoffmann, "Reduced storage and balancing needs in a fully renewable European power system with excess wind and solar power generation," *Renewable Energy*, vol. 36, no. 9, pp. 2515–2523, Sep. 2011.
[119] C. Budischak, D. Sewell, H. Thomson, L. Mach, D. E. Veron, and W. Kempton, "Cost-minimized combinations of wind power, solar power and electrochemical storage, powering the grid up to 99.9% of the time," *Journal of Power Sources*, vol. 225, pp. 60–74, Mar. 2013.
[120] M. G. Rasmussen, G. B. Andresen, and M. Greiner, "Storage and balancing synergies in a fully or highly renewable pan-European power system," *Energy Policy*, vol. 51, pp. 642–651, Dec. 2012.

Acta Universitatis Upsaliensis

Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology 1428

Editor: The Dean of the Faculty of Science and Technology

A doctoral dissertation from the Faculty of Science and Technology, Uppsala University, is usually a summary of a number of papers. A few copies of the complete dissertation are kept at major Swedish research libraries, while the summary alone is distributed internationally through the series Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology. (Prior to January, 2005, the series was published under the title "Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology".)

Distribution: publications.uu.se
urn:nbn:se:uu:diva-302837

ACTA
UNIVERSITATIS
UPSALIENSIS
UPPSALA
2016