Fluctuations of offshore wind generation - Statistical modelling

P. Pinson*, L.E.A. Christensen, H. Madsen
Informatics and Mathematical Modelling, Technical University of Denmark, R. Petersens Plads, Lyngby, Denmark
*pp@imm.dtu.dk

P.E. Sørensen
Wind Energy Department, Risø National Laboratory, Roskilde, Denmark

M.H. Donovan, L.E. Jensen
Dong Energy, Fredericia & København, Denmark

Abstract

The magnitude of power fluctuations at large offshore wind farms has a significant impact on the control and management strategies of their power output. If focusing on the minute scale, one observes successive periods with smaller and larger power fluctuations. It seems that different regimes yield different behaviours of the wind power output. This paper concentrates on the statistical modelling of offshore power fluctuations, with particular emphasis on regime-switching models. More precisely, Self-Exciting Threshold AutoRegressive (SETAR), Smooth Transition AutoRegressive (STAR) and Markov-Switching AutoRegressive (MSAR) models are considered. The particularities of these models are presented, as well as methods for the estimation of their parameters. Simulation results are given for the case of the Horns Rev and Nysted offshore wind farms in Denmark, for time-series of power production averaged at a 1, 5, and 10-minute rate. The exercise consists in one-step ahead forecasting of these time-series with the various regime-switching models. It is shown that the MSAR model, for which the succession of regimes is represented by a hidden Markov chain, significantly outperforms the other models, for which the rules for the regime-switching are explicitly formulated.

Keywords: wind power, offshore, fluctuations, statistical modelling, regime-switching, control, forecasting

1 Introduction

FUTURE developments of wind power installations are more likely to take place offshore, owing to space availability, less problems with local population acceptance, and more steady winds. This is especially the case for countries that already experience a high wind power penetration onshore, as for instance Germany and Denmark. This latter country hosts the two largest offshore wind farms worldwide: Nysted and Horns Rev, whose nominal capacities are of 165.5 and 160 MW, respectively. An overview of offshore wind energy in Europe is given in [1].

Such large offshore wind farms concentrate a high wind power capacity at a single location. Onshore, the same level of installed capacity is usually spread over an area of significant size, which yields a smoothing of power fluctuations [2]. This spatial smoothing effect is hardly present offshore, and thus the magnitude of power fluctuations may reach very significant levels. Modelling the power fluctuations for the specific case of offshore wind farms is a current challenge [3], for better forecasting offshore wind generation, developing control strategies, or alternatively for simulating the combination of wind generation with storage. The present paper investigates the applicability and performance of some statistical models.

Operators of offshore wind farms often observe abrupt changes in power production. The fast variations can be related to the turbulent nature of the wind. They are smoothed out when considering the cumulative production for the wind farm, since turbines are spread over a pretty large area. In addition, when inspecting power production data averaged at a few-minute rate, one observes variations that are due to slower local atmospheric changes e.g. frontline passages and rain showers [4]. The example of a 10-day episode with wind power production at Horns Rev, consisting of 10-minute averages, is depicted in Fig. 1. These meteorological phenomena add complexity to the modelling of wind power production, which is already non-linear and bounded owing to the characteristics of the wind-to-power conversion function. Such succession of periods with power fluctuations of lower and larger magnitudes calls for the use of regime-switching models. Here, it is explained how to apply the Self-Exciting Threshold AutoRegressive (SETAR) model, the Smooth Transition AutoRegressive (STAR) model, as well as the Markov-Switching AutoRegressive (MSAR) model for that purpose. Their performance are evaluated on a one-step ahead forecasting exercise, and compared to those of linear models, i.e. AutoRegressive Moving Average (ARMA) models. The available data consist in time-series of power production averaged at a 1, 5, and 10-minute rate, for the Horns Rev and Nysted wind farms.

2 From linear to regime-switching models

Generated wind power is considered hereafter as a stochastic process for which statistical models are set up in order to describe its temporal evolution. The notation y_t is used for denoting both the state of the stochastic process at time t and the measured value at that time. All the measured power values over the considered period are gathered in the time-series $\{y_t\}$, $t = 1, \ldots, T$, where T is the total number of successive observations. The set $\Omega_t = (y_1, y_2, \ldots, y_t)$ that contains all the observations up to time t, is referred to as the information set. Our framework is that of univariate time-series modelling, i.e. no explanatory variable is used.

The well-known linear ARMA model is briefly introduced and will be used as a benchmark. It serves as a basis

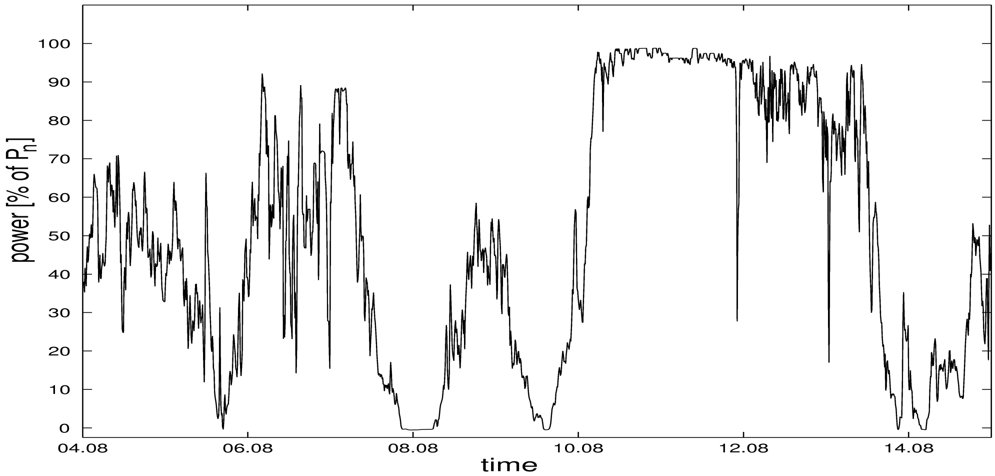

Figure 1: Wind power generated at Horns Rev over a 10-day episode in August 2005. Power values consist in 10-minute averages, normalized by the nominal power P_n of the wind farm (160 MW).

for constructing the regime-switching models. The term 'regime' originates from the assumption such that the considered stochastic process switches between a finite number of distinct (and most often linear) models. Denote by R the number of these regimes. The SETAR and STAR models are consequently presented, with focus given to the estimation of their parameters, and their use for time-series forecasting. For these two families of models, the switches from one regime to the other are governed by an observable process, i.e. by some function of lagged values of $\{y_t\}$.

2.1 The baseline ARMA model

The linear ARMA(p, q) model encompasses an autoregressive (AR) part of order p and a moving average (MA) part of order q

$$y_t = \theta_0 + \sum_{i=1}^{p} \theta_i y_{t-i} + \sum_{j=1}^{q} \phi_j \varepsilon_{t-j} + \varepsilon_t \tag{1}$$

where $\{\varepsilon_t\}$ is a white noise process, i.e. a purely random process with zero mean and variance ${\sigma_\varepsilon}^2 < \infty$. For the theory related to ARMA models we refer to e.g. [5]. Let us denote by Θ_a the parameter set, that is,

$$\Theta_a = (\theta_0, \ldots, \theta_p, \phi_1, \ldots, \phi_q, \sigma_\varepsilon)^\top \tag{2}$$

with $.^\top$ the transposition operator. Θ_a is obtained with Maximum Likelihood (ML) estimation with a Gaussian assumption on the distribution of residuals [5].

At time $t-1$, the set of parameters Θ_a can be used for calculating the one-step ahead point forecast $\hat{y}_{t|t-1}$ for the considered ARMA(p, q) process. This prediction corresponds to the conditional expectation of $\hat{y}_t$ given Θ_a and the information set Ω_{t-1}, and is readily given by

$$\hat{y}_{t|t-1} = \theta_0 + \sum_{i=1}^{p} \theta_i y_{t-i} + \sum_{j=1}^{q} \phi_j \varepsilon_{t-j} \tag{3}$$

ARMA models have already been applied for the modelling of wind power time-series. For instance, Milligan et al [6] have found them appropriate for producing 10-minute ahead forecasts of wind generation for onshore wind farms. In addition, Madsen [7] has shown that only little could be gained by applying more complex models i.e. bilinear or STAR models for such short horizons. Here, ARMA models are considered as a benchmark for comparison with the more advanced regime-switching models.

2.2 The SETAR model

A Self-Exciting Threshold AutoRegressive (SETAR) model is a piecewise linear model with an AR part for each of the R regimes, and for which the current regime is determined by a function of lagged values of the time-series [8]. This may yield abrupt switches from one regime to the other. Threshold values r_k, $k = 1, 2, \ldots, R-1$ define the intervals on which the various AR parts are active.

The SETAR$(R; p_1, p_2, \ldots, p_R)$ model is given by

$$y_t = \theta_0^{(m_t)} + \sum_{i=1}^{p_{m_t}} \theta_i^{(m_t)} y_{t-i} + \sigma_{m_t} \varepsilon_t \tag{4}$$

where for a given regime k, p_k and ${\sigma_k}^2$ denote the order of the AR model and the related variance of the noise sequence. $\{\varepsilon_t\}$ is a white noise process with unit variance, such that ε_t is independent of Ω_{t-1}. $\{m_t\}$ is the sequence of regimes, taking values in $\{1, 2, \ldots, R\}$, for which each m_t is defined by

$$m_t = \begin{cases} 1, & y_{t-d} \in \ ]-\infty; r_1] \quad \text{(regime 1)} \\ 2, & y_{t-d} \in \ ]r_1; r_2] \quad \text{(regime 2)} \\ \vdots & \vdots \\ R, & y_{t-d} \in \ ]r_{R-1}; \infty[\quad \text{(regime } R\text{)} \end{cases} \tag{5}$$

with d seen as a lag parameter.

The parameters of the SETAR model are estimated with the Minimum Mean Square Error (MMSE) estimation method. Let us write the parameter set Θ_e for the SETAR model as

$$\Theta_e = (\boldsymbol{\theta}, \mathbf{r}, \boldsymbol{\sigma})^\top \tag{6}$$

with

$$\begin{aligned}
\boldsymbol{\theta} &= (\theta_0^{(1)}, \ldots, \theta_{p_1}^{(1)}, \ldots, \theta_0^{(R)}, \ldots, \theta_{p_R}^{(R)})^\top & (7)\\
\mathbf{r} &= (r_1, r_2, \ldots, r_{R-1})^\top & (8)\\
\boldsymbol{\sigma} &= (\sigma_1, \sigma_2, \ldots, \sigma_R)^\top & (9)
\end{aligned}$$

that is, as the collection of the AR model coefficients, the vector of threshold values, and the vector of standard deviations of the noise sequence in each regime, respectively.

Assume that the number of regimes and the order of each AR part are known. Then, the objective function to be minimized over a dataset of length T is

$$S(\Theta_e) = \sum_{t=p_{\max}+1}^{T} (y_t - \hat{y}_{t|t-1})^2 \tag{10}$$

where $p_{\max} = \max(p_1, p_2, \ldots, p_R)$. $\hat{y}_{t|t-1}$ denotes the one-step ahead prediction, which can be readily obtained as the conditional expectation of y_t given Ω_{t-1} and the set of parameters Θ_e

$$\hat{y}_{t|t-1} = \theta_0^{(m_t)} + \sum_{i=1}^{p_{m_t}} \theta_i^{(m_t)} y_{t-i} \tag{11}$$

with the regime m_t at time t determined according to Eq. (5). The optimal parameter set $\hat{\Theta}_e$ is finally

$$\hat{\Theta}_e = \arg\min_{\Theta_e} S(\Theta_e) \tag{12}$$

The above minimization problem can reduce to a linear least squares problem for the estimation of the parameters of the AR models given the threshold values, by concentrating the sum of squares. In this case, the Weighted Least Squares (WLS) estimate $\hat{\boldsymbol{\theta}}$ of the AR parameters can be calculated as

$$\hat{\boldsymbol{\theta}}(\mathbf{r}) = \left(\tilde{\mathbf{x}}^\top \tilde{\mathbf{x}}\right)^{-1} \tilde{\mathbf{x}}\mathbf{y} \tag{13}$$

where $\tilde{\mathbf{x}}$ is a matrix for which every row contains past values of y_t for each regime multiplied with 0 or 1 depending on the regime sequence. For instance, if the process is in regime R at time t, the t^{th} row of $\tilde{\mathbf{x}}$ (denoted by $\mathbf{x}_t{}^\top$) is built as

$$\mathbf{x}_t{}^\top = \mathbf{x}_{t|m_t=R}{}^\top = (0, \ldots, 0, 1, y_{t-1}, \ldots, y_{t-p_R}) \tag{14}$$

From this WLS formulation of the AR parameter estimation, the objective function formulated in Eq. (10) simplifies to a function of the thresholds only, so that the optimal threshold values are found as

$$\hat{\mathbf{r}} = \arg\min_{\mathbf{r}} \sum_{t=p_{\max}+1}^{T} (y_t - \hat{\boldsymbol{\theta}}(\mathbf{r})^\top \mathbf{x}_t)^2 \tag{15}$$

and the corresponding AR parameter estimates $\hat{\boldsymbol{\theta}}(\mathbf{r})$ (in a MMSE sense) are finally computed with Eq. (13).

Since the above optimization problem might prove to have a lot of local minima, the initialization of the optimization process is crucial. Here, it has been initialized with different starting points spread over the set of possible $\mathbf{r}$, chosen after inspecting the data.

2.3 The STAR model

Often, abrupt changes between regimes are not satisfactory, even though separate regimes have been clearly identified. Smooth Transition AutoRegressive (STAR) models have been introduced in the literature in order to feature smooth (and controllable) transitions between regimes [9]. The Smooth Transition Bilinear (STBL) model, which belongs to the family of STAR models, has already been successfully applied for describing wind speed variations [7]. For one-step ahead forecasting of half-hourly averaged data, Madsen [7] has described its performance as slightly better than that of a simple AR(1) model. Focus is given here to the multiple-regime STAR (that we will, for convenience, refer to as STAR only), for which the value of the considered stochastic process $\{y_t\}$ at time t is given as a weighted average of several AR parts. The weights assigned to the AR parts are a function of lagged values of $\{y_t\}$. For a number of regimes R, with an AR model of order p_k in the k^{th} regime, the STAR$(R; p_1, p_2, \ldots, p_R)$ is given by

$$\begin{aligned}
y_t = \sum_{k=1}^{R-1} \Bigg(\left(\theta_0^{(k)} + \sum_{i=1}^{p_k} \theta_i^{(k)} y_{t-i} \right) \tilde{g}_k(z_t) \\
+ \left(\theta_0^{(k+1)} \sum_{j=1}^{p_{k+1}} \theta_j^{(k+1)} y_{t-j} \right) g_k(z_t) \Bigg) + \varepsilon_t
\end{aligned} \tag{16}$$

with

$$\tilde{g}_k(z_t) = 1 - g_k(z_t) \tag{17}$$

where $\{\varepsilon_t\}$ is a white noise process with variance $\sigma_\varepsilon{}^2$, and g_k is a smooth function that controls the transition between the k^{th} and $(k+1)^{\text{th}}$ regimes. $g_k(z)$ takes values in the unit interval. The regime variable z_t can be defined as a lagged value y_{t-d} of the stochastic process (d is then the lag parameter), or alternatively as an average of a set of lagged values.

The choice of the transition function depends on which type of behaviour is to be modelled. The two most popular transition functions are the exponential and logistic ones. The latter is chosen here, since it permits to more clearly separate the different regimes. The logistic function is a 2-parameter function defined as

$$g_k(z) = \left(1 + \exp\left(-\gamma_k(z - c_k)\right)\right)^{-1}, \quad \gamma_k > 0 \tag{18}$$

where γ_k is the slope parameter, which controls the transition speed between the regimes k and $k+1$, and c_k is the midpoint between these two regimes. Note that a STAR model with a logistic transition is equivalent to the SETAR model introduced above when $\gamma_k \to \infty$ with $z_t = y_{t-d}$.

The estimation method of the AR parameters for the STAR model is very similar to that described above for the case of SETAR models. Write

$$\Theta_s = (\boldsymbol{\theta}, \boldsymbol{\Gamma}, \mathbf{c}, \sigma_\varepsilon)^\top \tag{19}$$

the set of parameters, with

$$\begin{aligned}
\Gamma &= (\gamma_1, \gamma_2, \ldots, \gamma_{R-1})^\top & (20)\\
\mathbf{c} &= (c_1, c_2, \ldots, c_{R-1})^\top & (21)
\end{aligned}$$

which are the parameters of the transition functions, ${\sigma_\varepsilon}^2$ the variance of the white noise process, and with $\boldsymbol{\theta}$ the parameters of the AR models in each regime, as given by Eq. (7) for SETAR models. The MMSE estimate of Θ_s is obtained by minimizing an objective function $S(\Theta_s)$ that is equivalent to that of Eq. (10), but for which the one-step ahead prediction $\hat{y}_{t|t-1}$ is this time calculated with

$$\hat{y}_{t|t-1} = \sum_{k=1}^{R-1} \left(\left(\theta_0^{(k)} + \sum_{i=1}^{p_k} \theta_i^{(k)} y_{t-i} \right) \tilde{g}_k(z_t) + \left(\theta_0^{(k+1)} \sum_{j=1}^{p_{k+1}} \theta_j^{(k+1)} y_{t-j} \right) g_k(z_t) \right) \tag{22}$$

Then, assuming that R and the order of the AR part in each regime, as well as Γ and $\mathbf{c}$ are known, the MMSE estimate of $\boldsymbol{\theta}$ can be readily obtained from a WLS formulation

$$\hat{\boldsymbol{\theta}}(\Gamma, \mathbf{c}) = \left(\tilde{\mathbf{x}}^\top \tilde{\mathbf{x}}\right)^{-1} \tilde{\mathbf{x}}\mathbf{y} \tag{23}$$

where $\tilde{\mathbf{x}}$ is the weighted regression matrix, for which every row contains for each regime lagged values of $\{y_t\}$, weighted by the value of the transition function at the given time step. For instance, if considering a STAR$(2;1,1)$ model, with γ and c the parameters of the logistic transition function g, the t^{th} row of $\tilde{\mathbf{x}}$ is given by

$${\mathbf{x}_t}^\top = (\tilde{g}(z_t), y_{t-1}\tilde{g}(z_t), g(z_t), y_{t-1}g(z_t)) \tag{24}$$

From this WLS formulation, the MMSE estimate of Θ_s is obtained by minimizing a reduced form of the objective function, i.e.

$$\hat{\Theta}_s = \underset{(\Gamma,\mathbf{c})}{\arg\min} \sum_{t=1}^{T} (y_t - \hat{\boldsymbol{\theta}}(\Gamma, \mathbf{c})^\top \mathbf{x}_t)^2 \tag{25}$$

with an appropriate nonlinear optimizer. $\hat{\boldsymbol{\theta}}$ is consequently calculated with Eq. (23). Like for the case of SETAR models, the optimization process may be sensitive to the choice of initial values for Γ and $\mathbf{c}$, and may thus reach local optima. This optimization process is therefore initialized with a set of threshold values spread over the set of possible values, chosen after inspection of the data. In parallel, the initial Γ is chosen to be a unit vector.

3 A regime-switching model governed by a hidden Markov chain

The models described above rely on an observable process for determining the actual regime, which is determined as a function of past values of the process. Markov Switching AutoRegressive (MSAR) models propose an alternative to this observable regime-switching modelling, by allowing the switches to be governed by an unobservable process. It is assumed to be a Markov chain. A nice feature of such approach is that it permits to reflect the impact of some external factors on the behaviour of certain time-series [10]. Indeed, it has been found particularly suitable for modelling the temporal evolution of weather variables, such as daily rainfall occurrences [11] or wind fields [12, 13] especially because it manages to capture the influence of some complex meteorological features e.g. related to the motion of large meteorological structures. For the specific case of the fluctuations of offshore wind generation, our aim is to use this hidden Markov chain for describing meteorological features governing the regimes that cannot be determined from past values of measured power production only. MSAR models and the estimation of their parameters are briefly presented here. An extended description is available in [14].

3.1 Description of MSAR models

MSAR models resemble SETAR models in their formulation. If considering R regimes and AR models of orders $p_1, p_2, \ldots, p_R$ for each of these regimes, the corresponding MSAR$(R; p_1, \ldots, p_R)$ model is indeed given by

$$y_t = \theta_0^{(s_t)} + \sum_{i=1}^{p_{s_t}} \theta_i^{(s_t)} y_{t-i} + \sigma_{s_t} \varepsilon_t \tag{26}$$

where $\{\varepsilon_t\}$ is a white noise process with unit variance, ${\sigma_k}^2$ the variance of the noise sequence in the k^{th} regime, and $\{s_t\}$ the regime sequence. Even though the regime sequence for MSAR models is unobservable, it is assumed that $\{s_t\}$ follows a first order Markov chain on the finite space $\{1, \ldots, R\}$: the regime at time t is determined from the regime at time $t-1$ only, in a probabilistic way

$$P(s_t = j | s_{t-1} = i, s_{t-2}, \ldots, s_0) = P(s_t = j | s_{t-1} = i) \tag{27}$$

All the probabilities governing the switches from one regime to the other are gathered in the so-called transition matrix $\mathbf{P}$, for which the element p_{ij} represents the probability of being in regime j given that the previous regime was i, as formulated in Eq. (27). $\mathbf{P}$ is such that: *(i)* all the elements on a given row sum to 1 since the R regimes represent all regimes that can be reached at any time; *(ii)* all p_{ij} are positive in order to ensure ergodicity, which means that any regime can be reached eventually.

The set Θ_m of model parameters for MSAR models,

$$\Theta_m = (\boldsymbol{\theta}^{(1)}, \ldots, \boldsymbol{\theta}^{(R)}, \boldsymbol{\sigma}, \mathbf{P})^\top \tag{28}$$

gathers the parameters of the AR parts in each regime,

$$\boldsymbol{\theta}^{(j)} = (\theta_0^{(j)}, \theta_1^{(j)}, \ldots, \theta_{p_j}^{(j)})^\top, \;\; j = 1, \ldots, R \tag{29}$$

the standard deviation of the noise sequence in all regimes,

$$\boldsymbol{\sigma}^\top = (\sigma_1, \sigma_2, \ldots, \sigma_R)^\top \tag{30}$$

as well as the transition matrix $\mathbf{P}$.

As an illustration, a MSAR$(2;1,1)$ model is simulated here. The transition matrix $\mathbf{P}$ is such that

$$\mathbf{P} = \begin{pmatrix} 0.95 & 0.05 \\ 0.05 & 0.95 \end{pmatrix} \tag{31}$$

and the other model parameters (i.e. AR coefficients and variance in each regime) are

$$({\boldsymbol{\theta}^{(1)}}^\top, {\boldsymbol{\theta}^{(2)}}^\top, \boldsymbol{\sigma}^\top) = (1, .9, 5, .8, 0.8, 1.4) \tag{32}$$

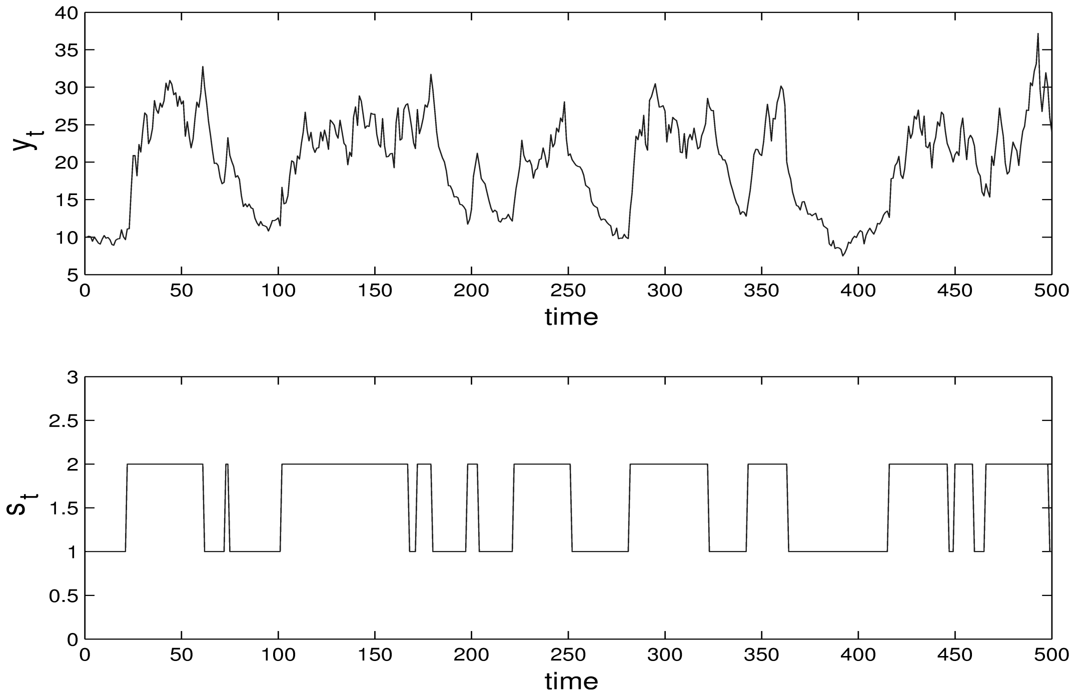

Figure 2: Simulation of a MSAR(2;1,1) process over a period of 500 time-steps. The transition matrix is given by Eq. (31), and the AR coefficients and variance in each regime by Eq. (32). Top: simulated process $\{y_t\}$. Bottom: regime sequence $\{s_t\}$.

The evolution of this MSAR process over a period of $T = 500$ time-steps is depicted in Fig. 2. The top part of the Figure shows the simulated process $\{y_t\}$, while the bottom part is related to the evolution of the regime sequence $\{s_t\}$. Owing to the choice of transition probabilities the switches between the two regimes are pretty rare.

3.2 Estimation

Estimating the parameters of MSAR models is more complicated than for the case of the regime-switching models with observable regime sequences. The method described in the following is based on the Expectation Maximization (EM) algorithm, which consists in an iterative method for maximizing the likelihood [15]. This two-step algorithm includes first an expectation step, for which the optimal inference of the regime sequence is determined, and a maximization step, where the parameters of the AR parts are updated by using the likelihood.

3.2.1 Optimal inference of regimes

A necessary assumption for determining the optimal inference of the regime sequence is that the number of regimes R, the order of the AR parts, as well as the set of parameters Θ_m are known. Even in this case, it is not possible to readily say in which regime the process belongs to for each observation. The solution to that problem is to consider a filtered probabilistic inference of the hidden regime sequence given the data [10]. Define the filtered probability $\hat{\xi}_{t|t}^{(j)}$ as the conditional probability of s_t being in regime j, given the information set Ω_t at time t and Θ_m, i.e.

$$\hat{\xi}_{t|t}^{(j)} = P(s_t = j|\Omega_t, \Theta_m) \tag{33}$$

These filtered probabilities for every regime can be arranged in a vector of filtered probabilities that we denote by $\hat{\boldsymbol{\xi}}_{t|t}$.

The filtered probabilistic inference allows one to iteratively calculate $\hat{\boldsymbol{\xi}}_{t|t}$ starting from $t = 1$, by drawing a simple relation between $\hat{\boldsymbol{\xi}}_{t|t}$ and $\hat{\boldsymbol{\xi}}_{t-1|t-1}$ given the observations up to time $t-1$ and the model parameters. Following [14], this relation writes

$$\hat{\boldsymbol{\xi}}_{t|t} = \frac{\hat{\boldsymbol{\xi}}_{t|t-1} \odot \boldsymbol{\eta}_t}{\mathbf{1}_R{}^\top(\hat{\boldsymbol{\xi}}_{t|t-1} \odot \boldsymbol{\eta}_t)} \tag{34}$$

where $\odot$ denotes the element-wise multiplication, $\mathbf{1}_R$ is a vector of ones of dimension R, and where $\hat{\boldsymbol{\xi}}_{t|t-1}$ is given by

$$\hat{\boldsymbol{\xi}}_{t|t-1} = \mathbf{P}^\top \hat{\boldsymbol{\xi}}_{t-1|t-1} \tag{35}$$

Finally, $\boldsymbol{\eta}_t$ is the vector that gathers at time t the conditional densities of y_t, given that the regime sequence is in such or such regime. From a Gaussian assumption on the noise sequence in each regime, $\varepsilon_t|s_t \sim \mathcal{N}(0, \sigma_{s_t}{}^2)$, the j^{th} element of $\boldsymbol{\eta}_t$ is

$$\eta_{t,j} = \frac{1}{\sqrt{2\pi}\sigma_j} \exp\left(-\frac{(y_t - \mathbf{x}_{t,j}{}^\top \boldsymbol{\theta}^{(j)})^2}{2\sigma_j{}^2}\right) \tag{36}$$

with $\mathbf{x}_{t,j}{}^\top = (1, y_{t-1}, \ldots, y_{t-p_j})$. It is also assumed that the conditional density of y_t only depends on the current regime. Eqs. (34) and (35) are often referred to as the Hamilton filter. They can be used in an iterative manner in order the calculate the vector of filtered probabilities for all observations.

Similarly, define the smoothed probability $\hat{\xi}_{t|T}^{(j)}$, i.e. the conditional probability of being in regime j at time t given the whole dataset Ω_T and the model parameters Θ_m

$$\hat{\xi}_{t|T}^{(j)} = P(s_t = j|\Omega_T, \Theta_m) \tag{37}$$

so that the sequence of smoothed probabilities related to the regime sequence can be calculated recursively with

$$\hat{\boldsymbol{\xi}}_{t|T} = \hat{\boldsymbol{\xi}}_{t|t} \odot \left(\mathbf{P} \left(\hat{\boldsymbol{\xi}}_{t+1|T} \oslash \mathbf{P}^\top \hat{\boldsymbol{\xi}}_{t|t} \right) \right) \tag{38}$$

with $\odot$ and $\oslash$ the element-wise multiplication and division, respectively [14].

3.2.2 Estimation of the model parameters with the EM algorithm

The EM algorithm is an iterative method that permits for MSAR models to estimate the transition matrix, the parameters of the AR parts and the variance in each regime, with the aim of maximizing the conditional log-likelihood $\ell(\Theta_m|\Omega_T)$ of the model parameters. The first step of the algorithm i.e. the expectation step, relates to the inference of the regime sequence (cf. above Paragraph). The second part, i.e. the maximization step, consists in the application of a set of update equations for the subsets of models parameters $\boldsymbol{\theta}^{(j)}$, $\boldsymbol{\sigma}$, and $\mathbf{P}$. The maximization of the likelihood is carried by successively applying expectation and maximization steps. Even though the conditional log-likelihood is not directly used in the method, it can be shown that the EM algorithm asymptotically maximizes that likelihood [16]. In addition, the evolution of its value can be monitored since it can be computed at each expectation step [14].

Each element p_{ij} of the transition matrix $\mathbf{P}$ can be estimated by visiting the sequence of smoothed probabilities and by applying the following equation:

$$\hat{p}_{ij} = \left(\sum_{t=p_{\max}+1}^{T} \hat{\xi}^{(j)}_{t|T} \hat{\xi}^{(i)}_{t-1|T} \right)^{-1} \sum_{t=p_{\max}+1}^{T} \hat{\xi}^{(j)}_{t|T} \tag{39}$$

where $p_{\max}$ is the maximum order of the AR models.

Consequently, the parameters $\boldsymbol{\theta}^{(j)}$ of the AR part related to regime j can be re-estimated by using a WLS formulation. It is indeed possible since the probabilities of being in such or such regime at any time t are known. They are given by the sequence of smoothed probabilities $\{\hat{\boldsymbol{\xi}}_{t|T}\}$. In the WLS formulation, each observation is weighted by the smoothed probability of being in that regime.

For that purpose, arrange the smoothed probabilities related to regime j in a weight matrix $\boldsymbol{\Sigma}_j$ (of dimension $T \times T$), for which the t^{th} element on the diagonal corresponds to the smoothed probability of regime j at time t, that is,

$$\boldsymbol{\Sigma}_j = \text{diag}\left(\hat{\xi}^{(j)}_{1|T}, \ldots, \hat{\xi}^{(j)}_{T|T} \right) \tag{40}$$

Then, for each regime j, an estimate of the AR parameters $\boldsymbol{\theta}^{(j)}$ can be computed by solving the usual equation for WLS estimation

$$\hat{\boldsymbol{\theta}}^{(j)} = (\tilde{\mathbf{x}}_j^\top \boldsymbol{\Sigma}_j \tilde{\mathbf{x}}_j)^{-1} \tilde{\mathbf{x}}_j^\top \boldsymbol{\Sigma}_j \mathbf{y}_j \tag{41}$$

where

$$\tilde{\mathbf{x}}_j = \begin{pmatrix} \mathbf{x}_{1,j}{}^\top \\ \vdots \\ \mathbf{x}_{T,j}{}^\top \end{pmatrix}, \quad \text{and} \quad \mathbf{y}_j = \begin{pmatrix} y_{p_j+1} \\ \vdots \\ y_T \end{pmatrix} \tag{42}$$

Finally, updating the variance $\sigma_j{}^2$ in regime j can be readily done with

$$\hat{\sigma}_j{}^2 = \frac{1}{\mu_j(T-(2+p_{\max}))} \sum_{t=1+p_{\max}}^{T} \left(y_t - \mathbf{x}_{t,j}{}^\top \hat{\boldsymbol{\theta}}^{(j)} \right)^2 \hat{\xi}^{(j)}_{t|T} \tag{43}$$

that is, by summing over t the squared residuals of the AR model related to regime j, weighted by the smoothed probability of being in regime j at time t, and normalizing by μ_j the sum of the smoothed probabilities for this regime.

The EM algorithm starts from an initial value for $\hat{\boldsymbol{\xi}}_{1|0}$. It is chosen to set each element of $\hat{\boldsymbol{\xi}}_{1|0}$ to equal probability, i.e. $\hat{\xi}^{(j)}_{1|0} = R^{-1}, \forall j$. An alternative would be to integrate the initial value estimation in the likelihood maximization problem [17]. Finally, the initial transition probabilities gathered in $\mathbf{P}$, as well as the initial variances and AR parameters, are derived after inspection of the dataset.

3.3 Forecasting with MSAR

At time $t-1$, producing a one-step ahead forecast with MSAR models consists in determining the conditional expectation of y_t given the information set Ω_{t-1} and the model parameters Θ_m. For that purpose, one first needs to predict the probabilities to be in such or such regime at time t with

$$\hat{\boldsymbol{\xi}}_{t|t-1} = \hat{\mathbf{P}}^\top \hat{\boldsymbol{\xi}}_{t-1|t-1} \tag{44}$$

where $\hat{\mathbf{P}}$ is the transition matrix estimated over the training set, and $\hat{\boldsymbol{\xi}}_{t-1|t-1}$ gathers the probabilities of being in such or such regime at time $t-1$ (cf. definition (34)). Then, the one-step ahead prediction for the stochastic process itself is calculated as

$$\hat{y}_{t|t-1} = E(y_t|\Omega_{t-1}, \hat{\Theta}_m) = \hat{\mathcal{A}}_t^\top \hat{\boldsymbol{\xi}}_{t|t-1} \tag{45}$$

i.e. as the weighted sum of the AR forecasts for every regime ($\hat{\mathcal{A}}_t$), the weights being given by the probabilities of being in these regimes.

4 Results from offshore case studies

The models presented above are used for describing the fluctuations of offshore wind generation on two real-world case studies. The exercise consists in one-step ahead forecasting of time-series of wind power production. The data for these two offshore wind farms are firstly described. Then, the configuration of the various models and the setup of estimation methods are given. Finally, a collection of results is shown and commented.

4.1 Case studies

The two offshore wind farms are Horns Rev (160MW) and Nysted (165.5MW), located in Denmark, off the west coast of Jutland and off the south cost of Zealand, respectively. The annual energy yield for each of these wind farms is around

600GWh. They are the two largest offshore wind farms worldwide today.

The raw power data consist in one-second measurements for each wind turbine. Following Sørensen et al [4], it has been chosen to model each wind farm as a single wind turbine, the production of which consists in the average of the power generated by all the available wind turbines. These turbines are of nominal capacity 2000 kW and 2300 kW for Horns Rev and Nysted, respectively. A sampling procedure has been developed for obtaining time-series of 1, 5, and 10-minute power averages. These sampling rates are selected so that the very fast fluctuations related to the turbulent nature of the wind disappear and reveal slower fluctuations at the minute scale. Because there may be some erroneous or suspicious data in the raw measurements, it has been considered that a minimum of 75% of the raw measurement within a time interval needed to be available so that the related average is seen as valid. At Horns Rev, the available raw data are from 16th February 2005 to 25th January 2006. And, for Nysted, these data have been gathered for the period ranging from 1st January to 30th September 2005.

The time-series have been splitted into learning and testing sets. The former serves for the fitting of statistical models while the latter allow us to appraise what the performance of these models may be in operational conditions. Sufficiently long periods without any invalid data are identified in order to define the necessary datasets. For Horns Rev, the training set relates to September 2005. The testing set is composed by 19 periods whose lengths are between 2 and 16 days, identified in the remaining of the whole dataset. Regarding Nysted, the training set corresponds to the period from the 15th February 2005 to the 9th March 2005, while the test set gathers 14 periods of length 6-27 days from the rest of the available data.

4.2 Models, estimation setup and evaluation criteria

The various time-series are modelled with the linear ARMA and regime-switching SETAR, STAR and MSAR models. The order of the AR and MA parts are chosen to vary between 1 and 5. This yields 25 competing ARMA models.

The optimal threshold values for SETAR and STAR models are determined from the nonlinear optimization procedures described in Paragraphs 2.2 and 2.3. The lag parameter d is chosen to be 1. We impose the number of regimes to be $R = 3$. Our choice for 3 regimes is motivated by the influence of the wind-to-power conversion function on the variance of wind generation: this variance is lower in the low and high power range, while it is much larger in the steep slope part of the power curve [18]. Thresholds are initialized by considering various combinations of lower and higher threshold values. The lower ones are picked in the set $\{200, 500, 800\}$ for both wind farms, while the higher ones are picked in the sets $\{1300, 1600, 1900\}$ and $\{1500, 1800, 2100\}$ for Horns Rev and Nysted, respectively. This yields 9 combinations of initial threshold values for each wind farm. For the particular case of STAR models, we fix the shape of the logistic functions by setting the slope parameter γ to 1. There are then 1125 competing models in each of the SETAR and STAR model families.

The initialization of the EM algorithm for the case of MSAR models consists in picking an initial transition matrix $\mathbf{P}$, as well as initial AR parts, by specifying their parameters and their variances, such that the resulting MSAR model is stationary. A stationary MSAR model is defined as a MSAR model whose AR part in each regime is stationary (cf. the definition of a stationary AR model given in [5]). The approach chosen here is to impose the transition matrix and the variances of the AR parts, while having the set of AR parameters varying. The initial $\mathbf{P}$ and $\boldsymbol{\sigma}$ are

$$\mathbf{P} = \begin{pmatrix} 0.8 & 0.1 & 0.1 \\ 0.1 & 0.8 & 0.1 \\ 0.1 & 0.1 & 0.8 \end{pmatrix}, \text{ and } \boldsymbol{\sigma} = \begin{pmatrix} 40 \\ 40 \\ 40 \end{pmatrix} \tag{46}$$

The initial parameters of the AR parts are chosen to take point of departure in the simple three regime MSAR$(3; 1, 1, 1)$ model with

$$\left(\boldsymbol{\theta}^{(1)\top}, \boldsymbol{\theta}^{(2)\top}, \boldsymbol{\theta}^{(3)\top}\right) = (1, 0.7, 50, 0.9, 100, 0.9) \tag{47}$$

Then, when increasing the order of one of the AR part, the new AR parameter to be initialized is given by a randomly chosen real number. The stationarity of the resulting MSAR model is verified. If this model is not stationary, another random number is drawn. This procedure is repeated until a stationary MSAR model is obtained. For each order of the MSAR model, we consider 10 different initial parameter sets yielding a stationary model. This raises the number of competing MSAR models to 1250.

Either with the ML estimation method and a Gaussian assumption on the residual distributions, or with the MMSE estimation method, the parameters of the models are determined with the aim of minimizing a quadratic error criterion. Therefore, in order to be consistent with the way parameters are estimated, models are also evaluated with a quadratic criterion on the testing set. More precisely, from the large panel of error measures available for evaluating wind power predictions [19], the Root Mean Square Error (RMSE) criterion is chosen.

4.3 Results and discussion

Table 1 lists the best models of each class — best in terms of a minimum RMSE on the testing set — for the time-series related to the Nysted wind farm. For instance for the 1-minute averaged data, the best of the 25 competing ARMA models has been found to be the ARMA(5,4). In addition, models are ranked from minimum to maximum RMSE. Table 1 also gives the characteristics of the optimal SETAR and STAR models, i.e. the thresholds that were determined from the optimization procedure. Note that for the 5 and 10-minute averaged data, the thresholds related to the lower regime for the SETAR models are very low (equal to 2.2 and 6 kW, respectively), thus isolating the no-production cases as a regime itself.

Whatever the sampling rate, the ARMA, SETAR and STAR models have a similar level of performance, while the

Table 1: Performance evaluation for the various models for Nysted. Results are for the 3 time-series averaged at different rates. The left column gives the optimal model of each class. The optimal threshold values for SETAR models (**r**) and STAR models (**c**) are also given. The models are ranked as a function of their RMSE on the testing set.

(a) 1-minute averaged data

Model	RMSE [kW]	**r** [kW]	**c** [kW]
MSAR(3;4,4,3)	13.3	–	–
STAR(3;5,5,5)	16.1	–	(920.9, 2096.6)
SETAR(3;4,4,4)	16.5	(203.6, 2006.3)	–
ARMA(5,4)	16.5	–	–

(b) 5-minute averaged data

Model	RMSE [kW]	**r** [kW]	**c** [kW]
MSAR(3;4,5,5)	35.5	–	–
STAR(3;5,5,5)	48.9	–	(827.3, 1638.7)
ARMA(4,5)	50.8	–	–
SETAR(3;1,3,3)	50.9	(2.2, 2149.8)	–

(c) 10-minute averaged data

Model	RMSE [kW]	**r** [kW]	**c** [kW]
MSAR(3;2,4,4)	60.6	–	–
STAR(3;5,5,5)	86.2	–	(579.4, 1545.4)
SETAR(3;3,5,5)	88.6	(6.0, 1595.4)	–
ARMA(5,1)	88.9	–	–

RMSE for the MSAR models is much lower. The improvement obtained with the Markov-switching models with respect to the three other types of models ranges from 19% to 32% depending on the sampling rate. Then, STAR models have an advantage versus the two others since the best STAR model is ranked second in all cases. The performance of the SETAR and ARMA models are a lot alike. It appears that considering separate regimes does not give any improvement against the classical linear models unless the switches between regimes are smoothed and controlled by some transition function. And, the hypothesis of some succession of regimes that could be captured with a first order Markov chain is validated by these results.

The testing set for Nysted is composed by 14 periods of different lengths and with different characteristics e.g. various mean production levels. The detail of the performance of the various models is shown in Fig. 3, which gives the RMSE of the models listed in the above Table for each period. There are only few periods for which the level of performance of the MSAR is worse than that of the other models. In general, the performance of all models is similar from one period to the other, and it does not seem that certain type of conditions would advantage such or such type of model. Also, by noticing that the curves for ARMA, SETAR and STAR models lie on top of each other whatever the period, one understands that modelling the regime-switching with a lagged value of measured wind power output does not yield a more dynamic modelling of the power fluctuations.

The same type of exercise is carried out for the Horns Rev case study. Table 2 gives the sorted list of the best models of each category for the three sampling rates, as well as their characteristics.[1] The thresholds of the SETAR model for the 5-minute averaged data, and those of the STAR model for the 10-minute sample data, are very close, showing that we almost converged towards two-regime models. The number of regimes has been imposed here, based on the knowledge of the effects of the non-linear and bounded power curve on power fluctuations. However, the number of regimes could also be considered as a model parameter to be optimized in the future, in order to see its influence on the resulting model performance.

Table 2: Performance evaluation for the various models for Horns Rev. Results are for the 3 time-series averaged at different rates. The left column gives the optimal model of each class. The optimal threshold values for SETAR models (**r**) and STAR models (**c**) are also given. The models are ranked as a function of their RMSE on the testing set.

(a) 1-minute averaged data

Model	RMSE [kW]	**r** [kW]	**c** [kW]
MSAR(3;3,2,5)	16.1	–	–
STAR(3;5,5,5)	20.1	–	(505.5, 1824.2)
SETAR(3;4,4,4)	20.4	(432.2, 1824.3)	–
ARMA(2,1)	20.6	–	–

(b) 5-minute averaged data

Model	RMSE [kW]	**r** [kW]	**c** [kW]
MSAR(3;3,1,5)	45.0	–	–
STAR(3;5,5,5)	63.3	–	(892.9, 1673.0)
SETAR(3;3,2,3)	65.1	(744.4, 760.6)	–
ARMA(2,2)	65.3	–	–

(c) 10-minute averaged data

Model	RMSE [kW]	**r** [kW]	**c** [kW]
MSAR(3;3,2,4)	68.1	–	–
STAR(3;5,4,5)	96.9	–	(705.0, 779.9)
SETAR(3;3,3,1)	99.8	(240.2, 2300)	–
ARMA(5,2)	99.9	–	–

Again, the STAR models have a slight advantage against the SETAR models, and these latter ones are also slightly better than linear ARMA models. But, they are significantly outperformed by the MSAR models, whatever the sampling rate. Indeed, the improvement proposed by this class of models with respect to the others ranges between 20 and 32%. This confirms once again the interest of considering a hidden Markov chain for modelling the regime-switching. In parallel, note that both for the Nysted and Horns Rev test cases, the average level of RMSE increases as the sampling rate gets larger. The persistent nature of wind generation makes that actual wind power output can be more easily modelled from recent power measures when the lead time is shorter. In addition, the average level of RMSE is significantly larger for Horns Rev than for Nysted, and this whatever the sampling rate. Since the estimated models are globally unbiased, this reveals that the variance of the model residuals is higher for the former wind farm, and hence that the random part of the fluctuations have a larger magnitude. This is certainly due to a more turbulent wind at Horns Rev.

The performance of the various models listed in Table 2 are detailed in Fig. 4 for the 19 periods composing the Horns Rev testing set. Those performance are more variable than for the previous test case. However, MSAR models are still significantly better than the other models for almost all periods, except for periods number 15 and 17. A particularity

[1] The upper threshold for the optimal SETAR model has converged to the nominal power value, indicating that this optimal model is indeed a SETAR(2;3,3) model.

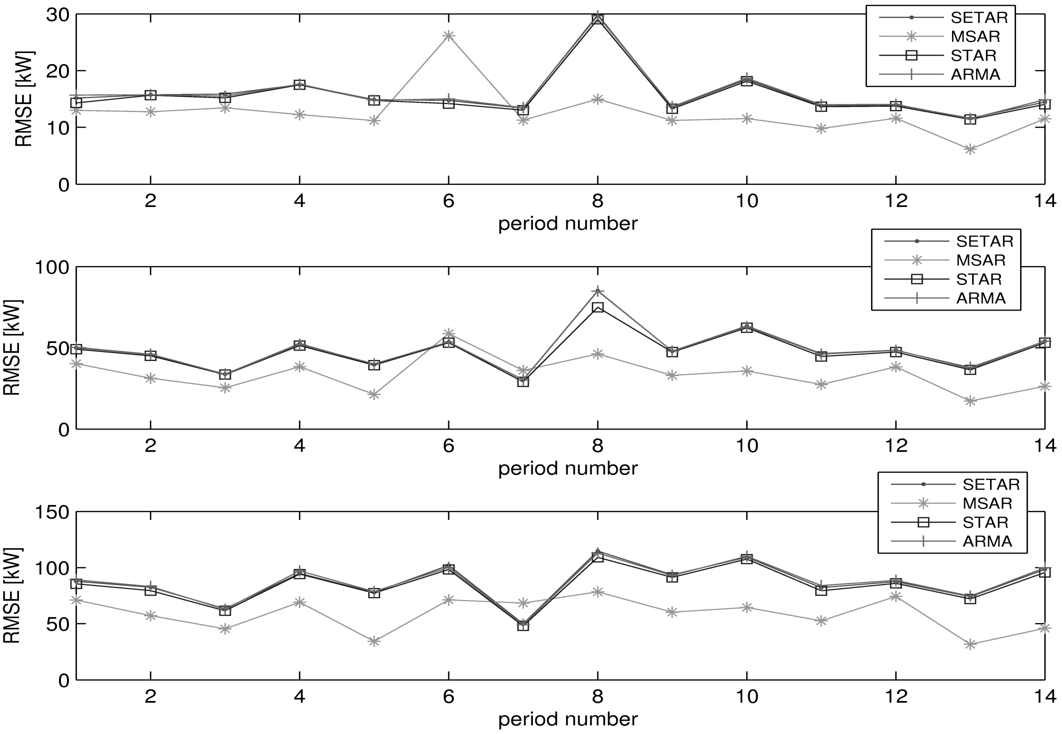

Figure 3: The RMSE on all test data sets from Nysted for each model. Top: 1 minute. Middle: 5 minute. Bottom: 10 minute.

of these two periods is that they consist in fast successions of drops and increases of wind power output. One may think that in these specific periods the SETAR and STAR models may be more appropriate since they have different AR parts depending on the level of power output, while it is not the case for MSAR and ARMA models. Though, since SETAR and STAR models do not exhibit a more significant improvement with respect to ARMA models for these two periods, this reveals that the regime-switching based on lagged values of power output does not have a higher value in these situations. Therefore, the poorer performance of MSAR models over periods 15 and 17 may simply be explained by the fact that the probabilistic inference of the regime-sequence was not very representative over these periods, owing to some more seldom meteorological phenomena.

5 Concluding remarks

Particular attention has to be given to the modelling of the fluctuations of offshore wind generation, since dedicated models are needed for enhancing the existing control and energy management strategies at offshore wind parks. This issue has been addressed by applying chosen statistical models. The choice for regime-switching approaches has been motivated by the succession of periods with fluctuations of lower and larger magnitudes that can be easily noticed when inspecting time-series of offshore wind power production averaged at a minute rate.

Two different types of regime-switching approaches have been applied. On the one hand, SETAR and STAR models rely on explicit rules for determining what the current regime is. It is in practice given by some function of past values of measured power. On the other hand, MSAR models are based on the idea that the regime-switching is governed by a hidden Markov process. This, from a theoretical point of view, may allow one to capture some complex influence of meteorological conditions on the wind power fluctuations. For verifying this (a priori) nice feature of MSAR models, they have been compared to SETAR, STAR and ARMA models on a one-step ahead forecasting exercise, with the aim of minimizing a quadratic error criterion. In all cases, it has been found that MSAR models significantly outperform the other ones: the error reduction ranges between 19 and 32% depending on the test case and the sampling rate. The gain of applying SETAR or STAR models instead of simple linear ARMA models does exist, but is relatively small. MSAR models indeed manage to capture the influence of some complex meteorological features on the power fluctuations. It will be of particular interest to study the relation between the temporal evolution of some meteorological variables and the regime sequences of MSAR models in order to determine which of these variables have a direct impact on the magnitude of power fluctuations. Integrating this knowledge in existing forecasting methods will permit to significantly increase their skill for the specific case of very short-term prediction (from some minutes to few hours) at offshore sites.

The results obtained with Markov-switching approaches encourage further investigation. The AR part in each regime could be extended to Generalized AutoRegressive with Conditional Heteroskedasticity (GARCH) models. Alternatively, we may propose to use AR models whose parameters are conditional to the level of the predictand. In such a case, the Gaussian assumption must be rethought, since conditional distributions of wind generation given the level of power output are not Gaussian. If a parametric assumption is to be

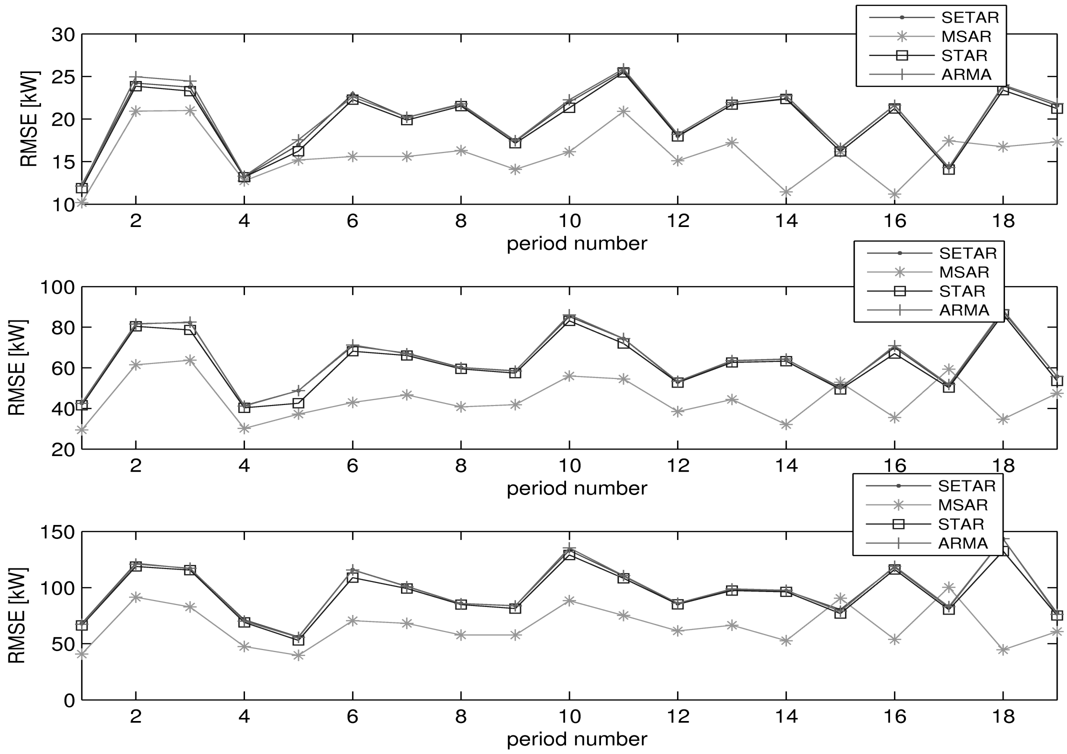

Figure 4: The RMSE on all test data sets from Horns Rev for each model. Top: 1 minute. Middle: 5 minute. Bottom: 10 minute.

made, a β-distribution assumption is much more suitable. The development and application of conditional β-MSAR models will be the focus of further research works. Finally, as wind generation is a non-stationary process it would be appropriate to have models with time-varying parameters.

Broader perspectives regarding follow-up studies include the development of stochastic models for simulating the interaction of offshore wind generation with conventional generation or storage, used as a backup for smoothing the fast power fluctuations at offshore wind farms. Better control strategies will result from the application of these models, which will significantly reduce the potential large costs induced by unwanted large power fluctuations.

Acknowledgements

The results originate from the project ‘Power Fluctuations in Large Offshore Wind Farms’ sponsored by the Danish PSO fund (PSO 105622 / FU 4104), which is hereby greatly acknowledged. Energi E2 and Elsam (now Dong Energy) are also acknowledged for providing the wind power data.

References

[1] IEA, 2005. Offshore Wind Experiences. IEA Publications, Paris, France.

[2] Focken, U., Lange, M., Monnich, M., Waldl, H.-P., Beyer, H.-G., Luig, A., 2002. Short-term prediction of the aggregated power output of wind farms - A statistical analysis of the reduction of the prediction error by spatial smoothing effects. J. Wind Eng. Ind. Aerod. 90: 231-246.

[3] Hendersen, A.R., Morgan, C., Smith, B., Sørensen, H.C., Barthelmie, R.J., Boesmans, B., 2003. Offshore wind energy in Europe - A review of the state-of-the-art. Wind Energ. 6: 35-52.

[4] Sørensen, P., Cutululis, N.A., Hjerrild, J., Jensen, L.E., Donovan, M.H., Christensen, L.E.A., Madsen, H., Vigueras-Rodríguez A., 2006. Power fluctuations from large offshore wind farms. Nordic Wind Power Conference Proceedings.

[5] Madsen, H., 2004. Time Series Analysis (2nd edition). Technical University of Denmark: Lyngby (ISBN 87-643-00098-6).

[6] Milligan, M., Swartz, M., Wan, Y., 2003. Statistical wind power forecasting models. Results for U.S. wind farms. National Renewable Energy Laboratory, Golden (Colorado).

[7] Madsen, H., 1996. Models and Methods for Wind Power Forecasting. Eltra: Skærbæk (ISBN 87-87090-29-5).

[8] Tong, H., 1990. Non-Linear Time Series - A Dynamical Approach (1st edition). Oxford University Press: Oxford.

[9] Chan, K.S., Tong, H., 1986. On estimating thresholds in autoregressive models. J. Time Ser. Anal. 7: 178-190.

[10] Hamilton, J.D., 1989. A new approach to the economic analysis of nonstationary time-series and business cycles. Econometrica 57: 357-384.

[11] Robertson, A.W., Kirshner, S., Smyth, P., 2003. Hidden Markov models for modeling daily rainfall occurence over Brazil. Report UCI-ICS-03-27, Information and Computer Sciences, University of California, Irvine (California).

[12] Ailliot, P., Monbet, V., 2006. Markov switching autoregressive models for wind time series. J. Stat. Plan. Infer. (submitted).

[13] Ailliot, P., Monbet, V., Prevosto, M., 2006. An autoregressive model with time-varying coefficients for wind fields. Environmetrics 19: 107-117.

[14] Pinson, P., Christensen, L.E.A., Madsen, H., Sørensen, P., Donovan, M.H., Jensen, L.E., 2006. Regime-switching modelling of the fluctuations of offshore wind generation. J. Wind Eng. Ind. Aerod. (submitted).

[15] Dempster, A.P., Laird, N.M., Rubin, D., 1977. Maximum likelihood from incomplete data via the EM algorithm. J. Roy. Stat. Soc. B 39: 1-38.

[16] Bishop, C.M., 1995. Neural Networks for Pattern Recognition (1st edition). Oxford University Press: Oxford.

[17] Hamilton, J.D., 1994. Time Series Analysis. Princeton University Press: Princeton.

[18] Pinson, P., 2006. Estimation of the Uncertainty in Wind Power Forecasting. Ph.D. dissertation, Ecole des Mines de Paris, Paris, France.

[19] Madsen, H., Pinson, P., Nielsen, T.S., Nielsen, H.Aa., Kariniotakis, G., 2005. Standardizing the performance evaluation of short-term wind power prediction models. Wind Eng. 29: 475-489.