Energy storage sizing for wind power: impact of the autocorrelation of day-ahead forecast errors

Pierre Haessig, Bernard Multon, Hamid Ben Ahmed, Stéphane Lascaud, Pascal Bondon

▶ **To cite this version:**

HAL Id: hal-00863901

https://hal.archives-ouvertes.fr/hal-00863901

Submitted on 19 Sep 2013

Energy storage sizing for wind power: impact of the autocorrelation of day-ahead forecast errors

Pierre Haessig, Bernard Multon, Hamid Ben Ahmed[*],
Stéphane Lascaud[†], and Pascal Bondon[‡]

Abstract

Availability of day-ahead production forecast is an important step towards better dispatchability of wind power production. However, the stochastic nature of forecast errors prevents a wind farm operator from holding a *firm production commitment*. In order to mitigate the deviation from the commitment, an energy storage system connected to the wind farm is considered. One statistical characteristic of day-ahead forecast errors has a major impact on storage performance: errors are significantly *correlated along several hours*. We thus use a data-fitted autoregressive model that captures this correlation to *quantify the impact* of correlation on storage sizing.

With a Monte Carlo approach, we study the behavior and the performance of an energy storage system (ESS) using the autoregressive model as an input. The ability of the storage system to meet a production commitment is statistically assessed for a range of capacities, using a mean absolute deviation criterion. By parametrically varying the correlation level, we show that disregarding correlation can lead to an underestimation of a storage capacity by an order of magnitude. Finally, we compare the results obtained from the model and from field data to validate the model.

Nomenclature

P_{pre} Forecasted wind power production (day-ahead)

P_{prod} Actual wind power production

P^*_{grid} Power commitment to the grid (day-ahead)

P_{grid} Power actually delivered to the grid

P_{dev} Deviation from the commitment

E_{sto} Energy stored in the energy storage system (ESS)

E_{rated} Capacity of the ESS

SoE State of Energy of the ESS

P^*_{sto} Ideal storage power request for the ESS

P_{sto} Power actually absorbed by the ESS

σ_P Standard deviation of the forecast error

[*]P. Haessig, B. Multon, H. Ben Ahmed are with SATIE, ENS Cachan Bretagne, CNRS, UEB, av. Robert Schuman, 35170 Bruz, France (email pierre.haessig@bretagne.ens-cachan.fr)
[†]S. Lascaud is with the LME department of EDF R&D, Ecuelles, France
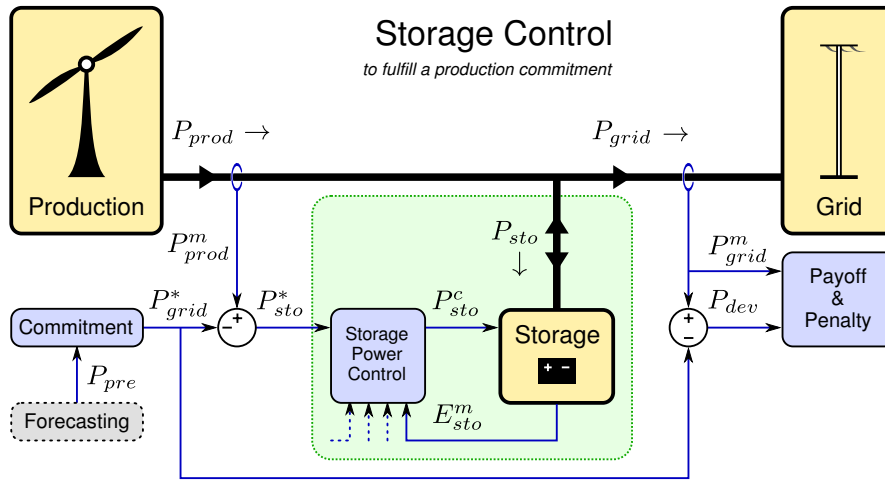[‡]P. Bondon is with L2S CNRS Supélec, Gif-sur-Yvette, France

Figure 1: Wind-storage system controlled to fulfill a day-ahead commitment

1 Introduction

In small islands with no interconnection with the power grid of a mainland, electricity production is often based on fossil fuels that are imported at a high cost. In such a context of expensive bulk electricity, wind turbines can be a profitable solution for electric power generation, wherever there is a strong enough wind power potential. Reducing CO_2 emissions induced by electricity generation is an additional motivation for a continuing growth of wind power in these islands.

However, the lack of grid interconnection makes the power systems of small islands rather weak. The *stochastic nature* of wind power production is a source of uncertainty in the grid energy balance which has a negative effect on grid stability. While wind power integration is already an issue for the large 300 GW European grid [15], maintaining grid stability becomes even more complicated for smaller grids because of their smaller inertia. Therefore, the growth of wind power, being labeled as *non-dispatchable*, is significantly hindered in some islands [11].

In addition, the wind induced energy unbalance is mainly mitigated by keeping extra spinning reserve, mostly fossil fuels based. Besides the investment cost issue of this additional generation capacity, the use of this spinning reserve generates CO_2 emissions which reduce the environmental benefit of wind power.

One first way to cut this costly need for spinning reserve is the use of *production forecasts* [3]. Because forecasting can significantly increase wind power dispatchability, this subject has received much academic and industrial attention for more than a decade, with numerous statistical and physical forecasting methods having been developed. A thorough review by Giebel et al. is available as part of the ANEMOS.plus European research project which focused on this subject [12].

Still, prediction is never perfect and the error magnitude, qualified in the sense of either Root Mean Square Error (RMSE) or Mean Absolute Error (MAE), is often about 10 % of the installed power capacity [12]. Prediction quality generally increases with the size of the region and the number of wind farms being considered. Small islands can thus suffer from larger prediction errors. For instance we report a 16 % RMS forecast error in the *real field data* used in this article, which comes from a wind farm in Guadeloupe, a French Caribbean island (see figure 2 for a week long data extract). A technical description of the farm can found in appendix 6.

Besides the forecast error magnitude, we want to show in this paper that a second parameter of major importance is the *correlation* along time of this error.

1.1 Energy Storage to Increase Wind Power Dispatchability

Energy storage is one way to mitigate the effect of prediction error. Compared to fossil fuel based spinning reserve, it is expected to be cleaner in terms of life cycle CO_2 emissions. Grid connected storage is an active research topic by itself [1], with specific attention being paid to wind-related storage [2].

In this article we study the wind-storage system of figure 1 which combines wind power generation with an energy storage system and day-ahead wind power forecasts to *fulfill a production commitment* to the grid operator. This system will be further described in the next sections.

A remarkable field demonstration of a wind-storage system is under construction on the island of El Hierro (Canarian archipelago) where a pumped-hydro storage coupled to wind generation was sized to reach a tremendous 80 % share of the yearly electricity production [8]. Pumped-hydro is very well adapted where a large *storage time constant* (the ratio of rated capacity to rated power) is required as this is the case on the island of El Hierro since the storage system is sized to supply the island during two full days without any wind, and as a further step, upgrading the autonomy to eight days is planned.

Electrochemical storage is another option. Depending on the technology, a wide interval of time constant, from minutes to many hours, is achievable. The electrochemical storage market potential is driven by both the availability of new technologies like Lithium-ion batteries and the increase in storage needs. Balancing the errors of wind generation forecasts is one of those. At the timescale of day-ahead wind power forecast, other technologies with smaller typical time constants like super-capacitors and flywheels are not appropriate. They are however well-adapted when smoothing out fast variations is at stake [5].

A country-wide analysis to assess how storage could help compensate for the difference between predicted and actual wind production was recently conducted by Hartmann and Dán for the Hungarian power grid [14]. Their *data-driven simulation tool* benefits from the availability of both production and forecast data for a long enough period (one full year). Unfortunately, such field data is not always easily available, especially prediction data because, at the moment, forecasting tools are not widespread in operational wind farms. In our case, field production/prediction data was only made available under a non-disclosure agreement which prevents data sharing which would be useful for creating an open benchmark for researchers.

One consequence of forecast data being generally unavailable is that most academic literature on the subject of wind-related storage follows what can be called a *stochastic approach*. This approach generally involves the use real production data (generally widely available) while the corresponding forecast data is *synthesized* from the production by adding a given amount of synthetic random error, thus the name "stochastic". We show in the following that this random error must be generated with care to *avoid serious underestimation* of the need for storage capacity.

1.2 Stochastic Approach: the Pitfall of Error Independence

The analysis of Teleke et al. [26] is one example of a combined wind-storage dispatching study using a stochastic approach. In their work, an hourly power dispatch set point was generated by "adding 10 % noise" to the actual hourly production average. More recently, Dicorato et al. [10], studying the market operation of a combined wind-storage system needed power forecasts to compute the day-ahead market bid. Inspired by Teleke, they generated synthetic *day-ahead* production forecasts at a 5 minutes timestep by adding a random error. They used a "linear" error distribution with standard deviation set to 10 % of the average power of that day. In a slight variation, Yao et al. [29] used a normally distributed error to perturb their measured wind speed data and error standard deviation was set to 10 % of their turbine rated wind speed. Adding the error to wind speed data instead of power was probably inspired by Lange's observations [18] that wind speed forecast errors exhibit more normality than power forecast errors.

In all these papers, as in others, forecast errors between successive instants were assumed to be *independent* random variables (error is a "white noise") even though this hypothesis was *not explicitly* stated. When forecast is used to make a single-timestep-ahead dispatch [26], this error independence holds true because the forecasting system can be updated with new information at each step. However, it doesn't hold true when it comes to forecasting blocks of upcoming timesteps, most common case being day-ahead dispatch [10]. Field data from many sources show that the hourly errors of a day-ahead forecast are *highly positively correlated*, and thus are not independent.

1.2.1 Motivations for incorporating autocorrelation

Taking into account correlation is important in all situations where wind forecasts are inputs of a *dynamical* system (a system with a memory the past). On a power grid, such systems are ubiquitous and include storage systems like hydro plants or batteries but also thermal units due to their commitment timing constraints

(e.g. minimum uptime). In this article, we quantitatively show that for a fixed amount of forecast standard error, the positive autocorrelation *severely increases* the need for storage capacity. Thus, error correlation should be taken into account for storage sizing in order to avoid what we may call the "error independence pitfall".

Because of its impact on power system analysis in general, forecast error correlation is already incorporated in some of the early studies on wind integration like Söder's pioneering 1994 report. In the dedicated chapter [24, p. 43-131] he describes with details how to model correlated wind forecast errors using an *autoregressive model*. Besides autocorrelation, his model can include the cross-correlations of errors between multiple production site. This work was subsequently expanded [6, 25] for the WILMAR wind integration project to generate scenarios (in particular scenario trees) which are an important input for certain dispatch algorithms. The autoregressive approach was also used for wind integration studies in Netherlands [28] or in California [9], the latter entering the details on their modeling method.

In relation to research on probabilistic forecasting tools, error interdependence was also underlined by Pinson et al. [21]. They were able to derive statistically generated production scenarios that indeed integrate correlated forecast errors. These scenarios were later used to study a dynamic storage sizing [22]. Their modeling approach is based on a *time-adaptive* dependence model. Also, the conditional error distribution is built with adaptive quantile regression so that it adapts to both the production level and the weather situation. On one hand, such a model has the advantage to be flexible enough to be applied to variety of geographical situations and to adapt to varying weather conditions. On the other hand, the large number of coefficients in this model requires the availability of an even larger amount of training data for model fitting. Therefore, there may be space for a simpler, approximate model that can be built on a limited amount of data and *still give useful trends for storage sizing*.

1.2.2 Motivations for a simple autocorrelation model

The analysis conducted in this article is based on a simple parametric model which incorporates error correlation, namely a *first order autoregressive process* [7] quite similar to Söder's [24]. Shortened "AR(1)", this model is detailed in section 3. It has only two parameters : one for the error magnitude and one that shapes an exponentially decaying autocorrelation. While its simplicity may prevent this model from properly representing *any* type of forecast error, it does appropriately fit the field data we have. Most importantly, despite some limitations (section 3.4), the model can *reproduce the performance* of an Energy Storage System fed with a field dataset for validation (section 4.5). This means it captures the most important statistical characteristics of forecast error *for this application*. In the end, the simplicity of the model enables a fast *parametric sizing analysis* from which a system designer can obtain a *sensible direction* for the choice of a storage capacity, even if only limited field data is available. Given the high unit cost of an energy storage, such a preliminary analysis does not preclude further studies when more data can be collected.

1.3 Scope of this Article

In this article, we take the stochastic approach we explained above to study the behavior of an Energy Storage System (hereafter ESS) associated with wind power generation of figure 1. A dedicated Python program was written to simulate the stochastic model using the Monte Carlo method. We parametrically assess the ESS performance by varying the coefficients of the stochastic AR(1) input model. Performance is assessed using power quality criterions detailed below, namely quadratic and absolute mean value (similarly to Teleke's Performance Index [26]). The parametric variation enables the *quantitative assessment* of the impact of forecast error correlation on the storage performance. This is our main contribution and we hope that this study can serve as a "quantitative argument" for taking correlation into account in all future wind-storage studies. Eventually, we translate the parametric analysis into a preliminary *ESS sizing method* which allows a designer to find which ESS capacity can *match a chosen performance level*.

Energy storage being a wide and complex issue, a lot of underlying questions were kept out of the scope of this article despite their importance. For instance, the issue of *storage profitability* is not thoroughly addressed because it depends on several parameters such as ESS investment cost (technology-dependent), ESS replacement cost (related to aging, which is an issue by itself). Profitability also depends on the business model that may be enabled by country specific laws adopted to encourage storage adoption. One way

to price stored energy is to require day-ahead production schedules from wind energy producers while putting financial penalties on schedule deviations. This kind of rule is already implemented country-wide in Hungary [14] and is being proposed for new farms in French islands like Guadeloupe (see below). However, a commonly shared rule to compute deviation penalties is yet to be defined, unless relying on the price fixing of an electricity balancing market.

The issue of *storage losses* is also not addressed in this article. Losses are often treated by simply considering a storage round-trip efficiency smaller than unity. However, we believe that the amount of losses is strongly coupled to the ESS technology, its capacity and to the amount of power flowing in and out (especially true for charging/discharging losses). Therefore we think that using a simple constant efficiency coefficient would bring little value to our study while a more complex treatment would bring us too far. However, for a particular storage technology (grid-scale sodium-sulfur batteries), we build upon the present work to include the technology-specific losses and costs models [13].

Finally, we do not model the interaction of the wind farm under study with neighbor farms (multisites correlations) or with the consumption (wind-load correlations). Indeed these important interactions are kept outside the industrial context considered here (see next section) because the policy of the local grid operator is to deal with *independent* commitment agreements with each wind power producer.

2 Description of the Wind-Storage System under Study

2.1 Regulatory Context of Wind Power in French Islands

All French islands like Guadeloupe (Caribbean) are subject to a national regulation which requires that, at any time of the day, uncertain electricity generation (i.e. wind and photovoltaic) should *never exceed 30 %* of the overall electricity consumption. Any excess of uncertain power generation is curtailed by disconnecting producers on a *Last In, First Out* basis. This regulation would weaken the profitability of any new wind farms in Guadeloupe so it may partly explain the significant slow down in new commissioning since 2008 [11], despite a high wind power potential.

In order to overcome this 30 % limit, a request for proposal was launched in November 2010 by the French Electricity Regulation Commission (CRE). It calls for the installation of a new kind of wind generation system, that would be more reliable (that is, dispatchable) as well as provide some ancillary services to the grid. Selected projects were announced in February 2012, for a total rated power of 95 MW. Those projects are now to be built in the near future.

2.2 Description of the Production Commitment Mechanism

The increase in production reliability is mainly based on a mechanism of day-ahead production commitment implemented by the wind-storage system of figure 1. Each day, before 5pm, the wind production operator has to announce a *power production schedule* for each half-hour interval in the following day. This schedule $P^*_{grid}(t)$, based on a production forecast $P_{pre}(t)$, is taken by the grid system operator as a firm *production commitment*. The actual power delivered to the grid $P_{grid}(t)$ should be kept as close as possible to this commitment. This rule is, to some extent, similar to the country-wide Hungarian requirements for wind power [14].

On the next day, the actual wind power production $P_{prod}(t)$ won't exactly match the commitment $P^*_{grid}(t)$, since forecasting is never perfect. In order to stay close to the commitment, an *Energy Storage System* is used. The ESS absorbs a certain amount of power $P_{sto}(t)$ so that the power delivered to the grid can be adjusted to some extent:

$$P_{grid}(t) = P_{prod}(t) - P_{sto}(t) \tag{1}$$

These three power flows are represented on figure 1 with bold lines. In this study, we do not integrate the possibility to curtail the production, that is the on-demand reduction of wind power generation, (e.g. by acting on the pitch angle of the blades). Thus $P_{prod}(t)$ is considered as a *fixed input* of the study. Therefore, $P_{sto}(t)$ is the only *degree of freedom* in equation (1). It is a bidirectional power flow: ESS charges when $P_{sto} > 0$ whereas it discharges when $P_{sto} < 0$.

Figure 2: One week extract of production and day-ahead prediction data, with a timestep of 1 hour.

2.3 Deviation from the Commitment

In order to assess the performance of the storage scheme, the deviation from commitment is defined :

$$P_{dev}(t) = P_{grid}(t) - P_{grid}^*(t) \tag{2}$$

From this deviation signal, a performance index can be derived, for instance by taking the quadratic or absolute mean value. The question of the criterion choice is discussed in section 4.2 below.

In order to ensure that this deviation is zero, the storage should absorb the exact difference between production and commitment. Therefore the *ideal storage power request* is expressed:

$$P_{sto}^*(t) = P_{prod}(t) - P_{grid}^*(t) \tag{3}$$

This request is "ideal" in the sense that whenever the actually stored power matches this ideal request ($P_{sto} = P_{sto}^*$) then, combining (1) and (3), the commitment is perfectly respected: $P_{grid} = P_{grid}^*$ and therefore the deviation P_{dev} is zero. The computation of both P_{sto}^* and P_{dev} signals is illustrated on figure 1.

For this article the production commitment P_{grid}^* is chosen equal to the production forecast P_{pre} (e.g. no advanced bidding strategy is used). Therefore the ideal storage request P_{sto}^* defined by (3) falls back to $P_{prod} - P_{pre}$ which is the *forecast error*. Thus we will use "forecast error" and "ideal storage request" interchangeably. The objective of the storage system is to compensate for this error by "absorbing" it.

A week long extract of actual production and prediction data (normalized to the rated power of the farm) is represented on figure 2. The difference between the two curves (the forecast error) is thus the ideal storage request P_{sto}^*, represented on the upper panel of figure 5. The standard deviation of P_{sto}^* being 16 % of the rated power, the need for storage capacity for this farm is expected to be large. More quantitative data about the forecast error is given when AR(1) modeling is applied in section 3.

2.4 System Simplification for Simulation Purpose

The wind-storage system of figure 1 has two input variables, namely P_{prod} and P_{pre}. However, only their *difference*, P_{sto}^*, is an input of the storage system. We can thus model and simulate an ESS with P_{sto}^* as the single output, without the need of production and the forecast variables. Figure 3 shows this simplified system, where P_{sto}^*, the forecast error, is generated by a random time series model, thus following the stochastic approach we introduced in section 1.2.

Furthermore, the assessment of the storage performance with P_{dev} defined by (2) can be conducted without the need of P_{grid} and P_{grid}^* which can thus be also removed from the simulation. Indeed, both can be substituted using (1) and (3) to obtain an *equivalent definition* of the deviation from the commitment:

$$P_{dev}(t) = P_{sto}^*(t) - P_{sto}(t) \tag{4}$$

With this additional transformation, the system of figure 3 becomes purely single-input single-output which enables a very lightweight simulation code. Notice that this wouldn't be possible if we dropped the

Figure 3: Simplified representation of the ESS for simulation purpose which behavior is described by (10) in section 4.1

assumption that *a)* the production P_{prod} cannot be curtailed and that *b)* the commitment is taken equal to the forecast.

3 Stochastic Modeling of Forecast Errors

The ideal storage request P_{sto}^* is the input of the storage system of figure 3. We've seen that it is equal to the forecast error which has a stochastic behavior. We now study its statistical properties and introduce a stochastic process that captures these properties. This statistical study, and especially time series analysis, was conducted with the free software R [23].

Since this article is based on hourly field data, we use discrete-time models, with a timestep $\Delta_t = 1\,\text{h}$. Also, all power values starting from here will be *average values* over that time step, i.e. equivalent to energy values.

3.1 Description and Statistical Study of Field Data

We study field data from the 1.375 MW "Grand Maison" wind farm, Guadeloupe island (see Appendix 6).

A day-ahead production forecast system has been running there since 2010, along with a continuous recording of power production at a 10 minutes timestep. Power production data was aggregated by averaging to match the 1 hour timestep of the forecast. Out of 327 days of available data, there was a significant proportion of missing values and/or period of weak power production. We then selected an extract of 50 days between June 3rd and July 23rd 2010 where power production was well measured. The first week of production and prediction (normalized with respect to the rated power) is depicted on figure 2 while the corresponding forecast error P_{sto}^* and is depicted on the upper panel of figure 5.

The standard deviation (std) σ_P of this forecast error over 50 days is found to be about 220 kW, that is 16 % of the rated power of the farm, while the average error, being 8 kW, is not significantly different from zero. This reassuringly shows that the forecast is not biased.

3.2 Dependence along Time of Forecast Errors

A simple way to analyze the dependence along time of a stationary process $\{P(k)\}$ is the use of the auto-correlation function, ratio of autocovariance and variance of the time series:

$$\rho_P(h) = \frac{\text{Cov}\big(P(k), P(k+h)\big)}{\text{Var}\big(P(k)\big)}, \quad h \in \mathbb{Z} \tag{5}$$

The estimated autocorrelation (acf function in R) of forecast errors P_{sto}^* is plotted on figure 4. Autocorrelation is as large as 0.8 for a single hour time lag and then decays exponentially. It stays significantly different from zero for up to 12 hours of lag and then enters the dotted blue 95 % confidence interval around zero.

We chose to capture the shape of this correlation structure with a process within the *autoregressive moving average* (ARMA) family [7]. ARMA processes are parametric linear models which parameters can be easily estimated using the arima function from R.

Figure 4: Autocorrelation of day-ahead forecast errors from the Grand Maison wind farm, compared with the theoretical autocorrelation (8) of an AR(1) model for three different values of the correlation coefficient ϕ.

3.3 Autoregressive Model Definition and Fitting

A zero-mean stochastic process $\{P(k)\}$ is called an ARMA(p,q) process if it satisfies the difference equation

$$
\begin{aligned}
P(k) - \phi_1 P(k-1) - \cdots - \phi_p P(k-p) = \\
\varepsilon(k) + \theta_1 \varepsilon(k-1) + \cdots + \theta_q \varepsilon(t-q),
\end{aligned}
\tag{6}
$$

where $\{\varepsilon(k)\}$ is a sequence of independently and identically distributed (IID) zero-mean random variables. Polynomials $\phi(z) = 1 - \phi_1 z - \cdots - \phi_p z^p$ and $\theta(z) = 1 + \theta_1 z + \cdots + \theta_q z^q$ with real coefficients have no common zeros and neither $\phi(z)$ nor $\theta(z)$ has zeros in the closed unit disk $\{z \in \mathbb{C} : |z| \leq 1\}$. There is a unique stationary solution $\{P(k)\}$ to eq (6).

Fitting an ARMA model to a given time series consists in selecting appropriate orders (p,q) and estimating the parameters. This is a statistical model selection problem which can be solved using information criterions such as AIC or BIC [7][section 9.3].

For our data $P_{sto}^*(k)$, these criterions choose the orders $(1,0)$ which corresponds to the simple AR(1) model

$$
P_{sto}^*(k) = \phi P_{sto}^*(k-1) + \varepsilon(k)
\tag{7}
$$

and the estimated value of ϕ is 0.78 (using the Gaussian conditional maximum likelihood method) with a 95 % confidence interval being $[0.75, 0.82]$. This confidence interval would have been thinner if a longer time series was considered. The estimated standard deviation $\hat{\sigma}_P$ of $\{P_{sto}^*(k)\}$ is 0.16 pu, that is 220 kW. The analysis of the residuals showed that there are uncorrelated, as expected.

It is easy to show that (7) implies

$$
\rho_{P_{sto}^*}(h) = \phi^{|h|}, \quad h \in \mathbb{Z}
\tag{8}
$$

and therefore ϕ is the correlation coefficient of $P_{sto}^*(k)$ between two successive instants. Observe that the empirical autocorrelation of the data (figure 4) decays exponentially with the time lags h, which is in good agreement with (8). In the case of our data, h is expressed in hours.

On figure 5 we present a side by side comparison of forecast error P_{sto}^* from both field data (upper panel) and stochastic simulations from AR(1) models (2nd and 3rd panels) with $\phi = 0.78$ and $\phi = 0$ respectively.

One can visualize the effect of the correlation coefficient ϕ on the *roughness* of the stochastic trajectories (this roughness is quantitatively assessed by the autocorrelation). The model with $\phi = 0.78$ and field data indeed share a similar roughness, as opposed to the model with $\phi = 0$.

8

Figure 5: Comparison between forecast error P_{sto}^* from field observations and two stochastic simulations (AR(1) model with $\sigma_P = 0.16$ in both cases while $\phi = 0.78$ and $\phi = 0$, in the middle and bottom panels respectively)

3.4 Analysis of Extra Data Sets, Model limitations

Beyond modeling the forecast error from the Grand Maison wind farm considered so far, we also analyzed some openly available data sets. From MAVIR, the Hungarian grid operator, we collected the year 2011 data used by Hartmann and Dán [14] (rated wind power of about 300 MW in 2011). We also collected data from Bonneville Power Administration (BPA) which openly provides wind power forecasts starting June 2012 for its area (rated wind power of about 4.5 GW in 2012). We give more details on this valuable dataset in the Appendix.

All time series are aggregated if necessary to hourly averages. Then, we compute the difference between day-ahead forecasts and actual production. RMS forecast error is 13 % of the rated power for MAVIR, while it is 8.5 % for BPA.

For the error interdependence, we observe that:

1. day-ahead forecasts errors are always highly positively correlated,

2. this autocorrelation is exponentially decaying (approximately),

3. fitting an AR(1) model yields a ϕ coefficient of 0.87 for MAVIR, 0.85 for BPA. In addition, De Mello *et al.* [9] report an hourly correlation coefficient of 0.93 for Californian ISO data.

To a good extent, these results generalize what is observed with the Grand Maison farm, with slightly higher correlation though. However, deeper analysis underlines some limitations of the AR(1) model, which are in fact classical limitations of the ARMA class in general.

First, an ARMA model is stationary and cannot incorporate some daily patterns which can be observed in the autocorrelation of BPA data for example. Daily patterns can be related to the fact that day-ahead forecast series are not generated hour by hour, but rather once a day as *blocks of 24 hours*. One possible solution is to simulate blocks of errors as in the WILMAR project [6] which enables a reset of the error at the start of each new day. However, when day-ahead forecast must be delivered several hours in advance (7 hours according to French TSO rules, 14 hours in De Mello *et al.* [9]), the RMS error at each first hour of the day is not zero (see fig 5) and the correlation with the last hour of the previous day is still positive.

Second, it is particularly visible with BPA data that the magnitude of forecast errors is quite *correlated with the power production*: low wind conditions yield low errors while strong winds yield large errors ("gusts of errors"). An ARMA model cannot capture such a switching behavior, but this could be addressed by adding an underlying Hidden Markov chain. Indeed, Markov-switching autoregressive (MSAR) models have been used to capture the wind regime changes [4] by many authors since at least 2004. MSAR was also used to model short term variability of wind power [20].

Third, the marginal distribution of an ARMA is Gaussian while power forecast errors are generally not [18]. For example, BPA errors are quite Laplacian, with a *clustering* of values around zero. Despite the existence of "Laplacian AR" models, we believe that Markov-switching AR processes are more promising because the clustering around zero is probably related to the gusts of errors. Therefore MSAR models are a promising path to expand forecast error modeling since they can address both the stationarity and the marginal density issues.

This being said, each model is forged *for a purpose* and for us it is the evaluation of the performance of an energy storage set up to absorb the forecast error. Therefore, after presenting the storage model, we perform a quantitative model checking (section 4.5) using Tsay's parametric bootstrap approach [27]. It shows that storage simulations fed with either *field data* or AR(1) trajectories indeed yield comparable performance results (performance criterion for the storage is defined in section 4.2.2).

3.5 About Stochastic Simulation of Day-ahead Forecast Errors

Autocorrelation coefficient ϕ obtained from AR(1) fitting is found to be in the 0.8–0.9 range in different data sets. Therefore, using an AR(1) model with ϕ in this range can be a good starting point when *no or few real data is available*. It is certainly better than sampling independent errors (equivalent to $\phi = 0.0$) as seen in figure 5, but one should be aware of the quite steep increase of deviation criterions when ϕ rises from 0.8 to 0.9 (see fig. 7). Hence, collecting an at least minimal dataset of a few month is still recommended.

For simulation purpose, one also need to choose σ_P, the std of the forecast error which relates to the "quality" of the forecast (in the RMSE sense). To match a given σ_P, it is useful to recall that the variances of $\varepsilon(k)$ and $P_{sto}^*(k)$ are linked because of eq (6). Therefore, one can prove that the std of $\varepsilon(k)$ should be set to $\sigma_P \sqrt{1 - \phi^2}$.

The value of σ_P should be chosen with respect to the size and location of the wind farm(s) being considered since it was found to vary between cases (16 % for the single farm in Guadeloupe down to 8.5 % for the aggregated production in the BPA area). Literature provides guidelines for selecting a sensible RMSE value [12].

Finally, we should again warn potential users of a forecast error model that every model is to some extent *application specific*. Indeed, forecast error models and the related models of wind power scenarios [21], can serve as inputs for a wide variety of systems like multistage optimizers or, in our case, a storage system simulation. Depending on the application, different features of the scenarios (like ramps or extreme production levels) may be influential. Therefore, a significant part of the quality checking of stochastic production scenarios should be done with application specific criterions. For our storage simulation, we use parametric bootstrap (section 4.5). For other use cases, Pinson & Girard [19] describe an event-based verification approach which can be used to track generic production events like ramps.

4 Behavior of the Storage System

We have developed a Python tool to simulate the storage system of figure 3 using the AR model described by (7) to generate the stochastic input. This tools builds upon the open computing packages NumPy and SciPy [17] and the plotting package Matplotlib [16].

It enables Monte Carlo simulations which involve generating *many* time series samples from the AR model and thus compute many trajectories of the system variables like P_{dev}. These trajectories are used to *statistically assess* the performance of the system, according to criterion defined in section 4.2. This performance assessment could not be performed analytically because of the non-linear behavior of the storage (saturation when empty or full).

Figure 6: Stochastic simulation of an ESS, with two different correlation values for the autoregressive input model: $\phi = 0$ on the left while $\phi = 0.80$ on the right. In both cases, $\sigma_P = 0.16$ pu and $E_{rated}^{norm} = 15$

4.1 Modeling of the Storage System

It appears that the ESS is defined by three parameters: σ_P and ϕ for the input power and E_{rated}, its capacity. In addition to these numbers, an *energy management policy* must be chosen. For this article, we assume the use of a simple policy which is "store the ideal storage request P_{sto}^* whenever it is *feasible*", where feasibility means that the storing action should not break the two constraints satisfied by the stored energy E_{sto}:

$$0 \leqslant E_{sto} \text{ and } E_{sto} \leqslant E_{rated} \tag{9}$$

with $E_{sto} = 0$ and $E_{sto} = E_{rated}$ meaning that the storage is respectively *empty* or *full*. This energy management policy leads to the following dynamical model of the storage system:

$$E_{sto}(k+1) = f_{sat}\Big(E_{sto}(k) + P_{sto}^*(k+1)\Delta_t \Big) \tag{10}$$

with f_{sat}, a saturation function, being defined as

$$f_{sat}(E) = \begin{cases} 0 & \text{if } E < 0 \\ E & \text{if } 0 \leqslant E \leqslant E_{rated} \\ E_{rated} & \text{if } E > E_{rated} \end{cases} \tag{11}$$

Then, the power P_{sto} that is *actually* stored is

$$P_{sto}(k+1) = \frac{E_{sto}(k+1) - E_{sto}(k)}{\Delta_t} \tag{12}$$

and it matches the storage request (i.e. $P_{sto} = P_{sto}^*$) whenever the storage is *not saturated*, that is neither empty nor full. When this condition is met, the deviation from schedule P_{dev} defined by (4) is zero and therefore the wind operator perfectly fulfills his commitment. We are now to further describe the behavior of the storage governed by this policy.

4.2 Performance Assessment

4.2.1 Qualitative Assessment

In order to get a feeling of what is the storage behavior, we represent two different stochastic simulations on figure 6. The forecast error standard deviation σ_P is set to 0.16 pu as observed in Grand Maison. The storage capacity E_{rated} is set to $15\sigma_P\Delta_t = 3.3$ MWh, with this choice being detailed in the next section. For comparison purpose, the correlation coefficient ϕ takes two different values: $\phi = 0$ (no correlation) on the left pane while $\phi = 0.8$ (as observed in field data) on the right.

11

In the 1st sub-panel the State of Energy (SoE) of the storage is represented:

$$SoE = \frac{E_{sto}}{E_{rated}} \tag{13}$$

which is always kept between 0 and 1 thanks to the saturation defined in (11). One can see that SoE is rising when the storage request P_{sto}^* (2nd sub-panel) is positive since the storage should absorb an excess of power production with regard to the production commitment. Accordingly, SoE is falling when the storage request is negative because the storage should compensate a lack of wind power production.

As long as the storage doesn't reach saturation ($0 < SoE < 1$), the storage request P_{sto}^* and the actually stored power P_{sto} are superimposed and thus the deviation P_{dev} (3rd sub-panel) is zero. Otherwise, there is an excess of power delivered to the grid ($P_{dev} > 0$) when the battery is full, and a lack of power ($P_{dev} < 0$) when the storage is empty.

When comparing the two panes, the amount of deviation is *significantly bigger* when the input is *highly correlated* ($\phi = 0.8$). Indeed, this high input correlation means that the forecast error can stay positive or negative for several consecutive timesteps. Therefore, there is less compensation between positive and negative errors.

After this visual assessment, we are now to quantitatively assess this deviation.

4.2.2 Performance Criterion

For our quantitative assessment, we need to choose a *performance criterion* which describes how the power deviation P_{dev} is not always zero as it should ideally be. Since the input of the storage system P_{sto}^* is a random variable, so is P_{dev} which requires the use of the expectation operator $\mathbb{E}[\cdot]$ to obtain non-random values. In the context of Monte Carlo simulations, estimating the expectation falls back to computing an average over many realizations.

Several criterions could be chosen but taking a norm of the random variable $P_{dev}(k)$ is simple and common. For this study, we restrict ourselves to the mean absolute deviation (MAD) $\|P_{dev}\|_1 = \mathbb{E}[|P_{dev}(k)|]$ but a quadratic norm, $\|P_{dev}\|_2 = \mathbb{E}[|P_{dev}(k)|^2]^{1/2}$, could be used as well. It would give more weights to bigger deviations which may be suitable in some applications.

With simulations yielding a series $\{P_{dev}(n)\}$ of N observations, MAD is estimated by:

$$\|\hat{P}_{dev}\|_1 = \frac{1}{N} \sum_{n=1}^{N} |P_{dev}(n)| \tag{14}$$

If a wind operator must pay penalties for every kWh out of the commitment (with no tolerance range), this criterion is proportional to the *average penalty* paid by the operator.

For example, the penalty in Hungary is about 17 EUR/MWh [14], however, a $\pm 30\,\%$ tolerance range is applied. The Performance Index defined by Teleke [26] is similarly based on a MAD criterion with some tolerance.

In order to better quantify the improving effect of the storage, we normalize the criterion $\|P_{dev}\|_1$ with respect to its value $\|P_{dev}\|_1^0$ if there were no storage. This normalization factor can be easily computed. Indeed, in the absence of storage, the deviation P_{dev} defined by (4) is simply equal to the forecast error P_{sto}^* and since the latter is governed by the Gaussian process (7) we have:

$$\|P_{dev}\|_1^0 = \mathbb{E}[|P_{sto}^*|] = \sqrt{\frac{2}{\pi}}\sigma_P \approx 0.80\,\sigma_P \tag{15}$$

with the coefficient $\sqrt{\frac{2}{\pi}}$ being simply the natural ratio of MAD to standard deviation for a Gaussian random variable. The normalized deviation criterion $\|P_{dev}\|_1^{\text{norm}}$ is thus defined by

$$\|P_{dev}\|_1^{\text{norm}} = \frac{\|P_{dev}\|_1}{\sqrt{2/\pi}\sigma_P} \tag{16}$$

so that a value $\|P_{dev}\|_1^{\text{norm}} = 0.5$ for example would mean that the storage has reduced the effect of the forecast error by a factor of 2.

4.2.3 Statistical Estimation of the Performance

The MAD criterion (14) is an average of N observations of P_{dev} which is a random variable since the input P_{sto}^* is random. Therefore, the estimation must be conducted with some care to get good enough estimates. We detail here the most important aspects of the computation.

First, the number of observations N must be *big enough* so that the variance of the estimator is small enough thanks to the law of large numbers. Having small enough variance is particularly important for big values of storage capacity because the deviation $\|P_{dev}\|_1^{\mathrm{norm}}$ gets small. For range of values of this article, we found $N = 10^5$ suitable.

Also, for the law of large numbers to apply, the N observations $P_{dev}(n)$ in (14) should be *statistically independent*. This means it not acceptable to take N consecutive observations since one can clearly see on figure 6 that consecutive deviation values $P_{dev}(k)$ are not independent. We overcame this problem by running N independent simulations in parallel, thus ensuring complete independence between the observations.

Finally, the N observations should be *independent from the initial conditions* of the simulation, that is in our case the initial State of Energy $SoE(0)$. Therefore, the storage simulation must be run along *enough timesteps* so that the effect of the initial state fades away. Again, the bigger values of storage capacity are the more problematic cases, because the higher amount of stored energy implies a higher amount of memory from the initial state. Inspired by the diffusion properties of random walks, we derived an empirical rule to automatically tune the length K of our simulations. Indeed, the stored energy behaves like a random walk in the absence of saturation. We thus set K to the integral part of $(E_{rated}/(\sigma_P.\Delta_t))^2$ which quadratically increases with the storage capacity. In simulations performed for figure 7, with biggest capacity being $300\,\sigma_P.\Delta_t$, the maximum simulation length was then $K_{max} = 300^2 \approx 10^5$ points long. Also, our empirical studies showed that this choice for K was indeed big enough to reach steady state.

4.3 Parametric Study of the Performance

The performance of the storage, in terms of the MAD criterion $\|P_{dev}\|_1^{\mathrm{norm}}$, was computed for a range of storage capacities E_{rated} so that the results can be used as hints for performance driven storage sizing. Also, we ran the computation for a range of correlation coefficient ϕ (between 0 and 0.90) to explore its effect on the need for capacity.

When studying the parametric behavior of the storage system, one can notice that only the *relative value* of the capacity E_{rated} with respect to the magnitude of the forecast error σ_P has an impact of the system behavior. Therefore, we express our simulations results in terms of a *normalized capacity* which is a dimensionless quantity:

$$E_{rated}^{\mathrm{norm}} = \frac{E_{rated}}{\sigma_P.\Delta_t} \tag{17}$$

We chose 30 points for $E_{rated}^{\mathrm{norm}}$, logarithmically spaced between 1 and 300. For Grand Maison, this corresponds to a capacity E_{rated} between 0.22 and 66 MWh, since $\sigma_P.\Delta_t = 220\text{ kW} \times 1\text{ h} = 220\text{ kWh}$. We picked 10 values between 0 and 0.9 for the correlation coefficient ϕ.

For each 30×10 pair of values $(E_{rated}^{\mathrm{norm}}, \phi)$ on the grid of parameters, we estimated the deviation criterion $\|P_{dev}\|_1^{\mathrm{norm}}$, following the methodology described in section 4.2 using $N = 10^5$ independent vectorized simulations. Longest simulations involved $N \times K_{max} \approx 10^{10}$ samples and took about 12 minutes for a single pair of parameters. However, running several simulation instances in parallel on a multicore processor reduced the simulation time of all 300 grid points to a few hours.

Results are plotted on figure 7 as a set of curves, one for each value of ϕ. Also, we put the parameter $E_{rated}^{\mathrm{norm}}$ on the y-axis, while the computation output $\|P_{dev}\|_1^{\mathrm{norm}}$ was laid on the x-axis in order make this curve readable as a *storage sizing tool*, as demonstrated in the next section.

Looking at the set of curves, one can see how any increase in the correlation coefficient ϕ *severely raises*the need for storage capacity. The curve for $\phi = 0.80$ is highlighted as being about equal to the value estimated on our field data from Grand Maison.

Figure 7: Storage capacity requirements as a function of the MAD performance criterion, for different values of the autocorrelation coefficient ϕ

4.4 Application to a Preliminary Storage Sizing

As an example based on power ratings from the Grand Maison wind farm, we suppose that MAD criterion $\|P_{dev}\|_1$ should be reduced from 0.127 pu (its natural value without storage) down to 0.05 pu by a novel grid code requirement. This is a reduction to 39 % of the no-storage value, that is $\|P_{dev}\|_1^{\text{norm}} = 0.39$. Reading from figure 7 with $\phi = 0.8$, matching this level of performance requires a relative capacity of 15. Therefore, $E_{rated} = 15\sigma_P\Delta_t = 3.3$ MWh.

This is a large capacity for a wind farm of 1.375 MW rated power. However, had the correlation not been taken into account, which is equivalent to say $\phi = 0.0$, the capacity requirement would have been only $E_{rated} = 2\sigma_P\Delta_t$ that is 0.44 MWh, an *underestimation* by a factor of about 7! Since the investment cost of a storage is roughly proportional to its capacity, taking into account the autocorrelation of forecast errors is therefore critical for the correct budgeting a wind storage system.

4.5 Model Checking with Bootstrap

We've seen in section 3.4 that the AR(1) model can match several important statistical characteristics of forecast errors but some others, like the "gusts of errors", are not represented. Since the purpose of the model is to simulate the performance of a storage, we use Tsay's parametric bootstrap [27] approach to assess the model ability to match this objective.

We implement bootstrap by first running the storage model (10) with the P_{sto}^* input coming from the BPA dataset. The time series is 7175 hours long and is preferred over Grand Maison because the latter is shorter. The ESS normalized capacity is set to 15 as in the previous sizing example. The simulation outputs the $P_{dev}(k)$ time series and we compute the MAD criterion $\|P_{dev}\|_1^{\text{norm}}$ which is 50.8 % as pointed on the left of figure 8. This is more than the 39 % of the previous example because the correlation coefficient of BPA data is higher ($\hat{\phi} = 0.85$).

Then, using the data-fitted AR(1) model, we generate 1000 independent time series $P_{sto}^*(k)$ of the same length as the data. With the storage simulation, we hence collect 1000 samples of the $\|P_{dev}\|_1^{\text{norm}}$ criterion. The distribution of these bootstrapped samples is also represented on figure 8, with their mean being 46.8 %.

We can observe that there is a significant variability of the deviation criterion within the bootstrapped

14

Figure 8: Storage MAD performance simulation with BPA data against 1000 bootstrapped time series of the same length (7175 h). Bootstrap samples are generated with the AR(1) model, either with data-fitted correlation (ϕ=0.85) or no correlation for comparison. Storage normalized capacity E_{rated}^{norm} is 15 in all cases. On the right side, same data is plotted in the context of figure 7.

samples. Computing the quantiles, we get a 95 % confidence interval of 41.2 – 52.5. This variability of the criterion is here solely due to the *finite length* of the input samples (about 300 days) and leads to two comments.

First, the criterion obtained from the BPA sample (50.8) is within the 95 % confidence interval of the bootstrapped samples (41.2 – 52.5), This means that the data and the model are indeed *compatible*. For comparison, the bootstrapped samples with $\phi = 0$ are superimposed which shows that this model is clearly incompatible with the data (underestimation of the deviation criterion by a factor of 5).

Then, as a side result, the variability of the bootstrapped statistics gives a useful indication of the uncertainty which surrounds the data-estimated criterion: we obtain a *measure of accuracy*. On the right side of figure 8, the 95 % confidence intervals of the bootstrapped statistics are superimposed with the "capacity requirement curves" of figure 7 to put these measures of variability in context. Getting information about the accuracy of an estimate is in fact the main motivation of people using the bootstrap method. Indeed, because the $P_{dev}(k)$ series is strongly non Gaussian (see fig 6), the classical formula "$\pm 1.96\sigma$" to compute a confidence interval do not hold true. We can use the width of the bootstrapped samples as a substitute: $52.5 - 41.2 = 11.3$. This shows that a data-based storage simulation would need a longer data input (more than 300 days) to yield a precise enough criterion. This will be possible in the near future thanks to BPA providing new data every month.

5 Conclusion

In the context of wind power production, we use a simple autoregressive model to capture the stochastic behavior of day-ahead forecast errors, in particular the autocorrelation structure of these errors. This model fits well field data from several origins (a single farm, a whole country) and the correlation coefficient ϕ (between two consecutive hours) is always found in the 80–90 % range.

Based on this model as its input, we use a Monte Carlo simulation tool to analyze the behavior and the performance of an energy storage system. The ESS is used to mitigate the forecast errors in order to fulfill a day-ahead production commitment. We parametrically assess the storage performance using a Mean Absolute Deviation criterion. We compute this criterion as a function of the storage capacity for various values of the autocorrelation coefficient ϕ. These curves can be used as a tool for a preliminary storage sizing driven by the system performance. More important, we show that discarding the autocorrelation in such an analysis can *underestimate the storage capacity by an order of magnitude*.

For further studies, we see three directions to improve our stochastic storage simulation tool. The first one is the improvement of the storage model to include charging and discharging losses, thermal and aging effects, as this was done for a sodium-sulfur battery [13] but is yet to be done for a battery in the Lithium-ion family. Then, the simulation should accept more advanced control strategies of the ESS, including the choice of a production commitment that may be different from the forecasted production (bidding strategy). Third,

the investment, operational and replacement costs of the storage need to enter the model in order to assess the economic benefit of an ESS. This includes thinking on which model is the most appropriate to penalize commitment deviations and this is a wide issue.

Finally, we should underline the importance of the prediction quality, measured by its RMS magnitude σ_P. Indeed, the simplest way to decrease the deviation criterion $\|P_{dev}\|_1$ at the output of the system is to decrease the magnitude of the input σ_P. It can be shown that, if the relative deviation $\|P_{dev}\|_1^{\text{norm}}$ is small enough (< 10 %), any decrease of σ_P leads to a quadratic decrease of $\|P_{dev}\|_1$. On the other hand, the effect of increasing the storage size is only linear. Therefore, if a system designer must allocate resources between a better wind power production forecast system and a bigger storage system, the former investment may prove way more valuable.

6 Wind Farm data detailed description

Aérowatt data used in this article comes from the wind farm of "Grand Maison", located on the shoreline of Guadeloupe, one of the French Caribbean islands. Operated by Aérowatt, this 1.375 MW farm consists of five "Vergnet GEV-MP" 275 kW wind turbines with pitch regulation and two-speed induction generators. These wind turbines are specifically designed to withstand the frequent hurricanes of this region. To this end, a specific tower design enables the turbines to be fully laid on the ground in advance, in case of a weather alert. Wind power forecasts comes from a weather services provider contracted by Aérowatt.

BPA data comes from Bonneville Power Administration, the U.S. operator for the Pacific Northwest region. The rated wind power in the BPA Balancing Authority is about 4.5 GW in 2013. Wind data used for model validation comes from two different file sources :

- Wind Generation in the BPA Balancing Authority (`transmission.bpa.gov/Business/Operations/Wind/`). Yearly files at a 5 minutes timestep, starting 2007.

- Wind Power forecasting data (`www.bpa.gov/Projects/Initiatives/Wind/Pages/Wind-Power-Forecasting-Data.aspx`). Monthly files at a 1 hour timestep, starting June 2012. Forecast provided every hour, up to 72 hours ahead.

Those two sources were assembled and aggregated to obtain the forecast error at a 1 hour timestep, from June 2012 to March 2013. The forecast time series was extracted from the forecast file by extracting 24 hours blocks generated at 5pm and covering the next day from 0:00 to 23:00. This choice matches the constraints of the power commitment for new wind farms in French islands.

Acknowledgments

References

[1] Handbook of Energy Storage for Transmission or Distribution Applications . Tech. Rep. 1008703, EPRI, Dec. 2002.

[2] EPRI-DOE Handbook Supplement of Energy Storage for Grid Connected Wind Generation Applications . Tech. Rep. 1007189, EPRI, Dec. 2004.

[3] Ahlstrom, M., Blatchford, J., Davis, M., Duchesne, J., Edelson, D., Focken, U., Lew, D., Loutan, C., Maggio, D., Marquis, M., McMullen, M., Parks, K., Schuyler, K., Sharp, J., and Souder, D. Atmospheric Pressure. *IEEE Power Energy Mag. 9*, 6 (nov.-dec. 2011), 97–107.

[4] Ailliot, P., and Monbet, V. Markov-switching autoregressive models for wind time series. *Environmental Modelling & Software 30*, 0 (2012), 92–101.

[5] Aubry, J., Bydlowski, P., Multon, B., Ben Ahmed, H., and Borgarino, B. Energy Storage System Sizing for Smoothing Power Generation of Direct Wave Energy Converters. In *3rd International Conference on Ocean Energy* (2010).

[6] Boone, A. Simulation of Short-term Wind Speed Forecast Errors using a Multi-variate ARMA(1,1) Time-series Model. Master's thesis, Apr. 2005.

[7] Brockwell, P. J., and Davis, R. A. *Time Series: Theory and Methods*, second edition ed. Springer Series in Statistics. Springer, 1991.

[8] Daniélo, O. El Hierro, l'Île Électrique. *Systèmes Solaires, le Journal des Énergies Renouvelables*, 201 (Jan. 2011), 88–97.

[9] de Mello, P. E., Lu, N., and Makarov, Y. An optimized autoregressive forecast error generator for wind and load uncertainty study. *Wind Energy 14*, 8 (2011), 967–976.

[10] Dicorato, M., Forte, G., Pisani, M., and Trovato, M. Planning and Operating Combined Wind-Storage System in Electricity Market. *IEEE Trans. Sustain. Energy 3*, 2 (april 2012), 209–217.

[11] EDF SEI. Bilan Prévisionnel de l'Equilibre Offre/Demande d'électricité — Guadeloupe. Tech. rep., July 2011.

[12] Giebel, G., Brownsword, R., Kariniotakis, G. N., Denhard, M., and Draxl, C. The state-of-the-art in short-term prediction of wind power: A literature overview. Tech. rep., ANEMOS.plus, 2011.

[13] Haessig, P., Multon, B., Ben Ahmed, H., Lascaud, S., and Jamy, L. Aging-aware NaS battery model in a stochastic wind-storage simulation framework . In *PowerTech 2013 Conference, Grenoble, France* (June 2013).

[14] Hartmann, B., and Dán, A. Cooperation of a Grid-Connected Wind Farm and an Energy Storage Unit — Demonstration of a Simulation Tool. *IEEE Trans. Sustain. Energy 3*, 1 (Jan. 2012), 49–56.

[15] Holttinen, H., Orths, A., Eriksen, P., Hidalgo, J., Estanqueiro, A., Groome, F., Coughlan, Y., Neumann, H., Lange, B., Hulle, F., and Dudurych, I. Currents of Change. *IEEE Power Energy Mag. 9*, 6 (nov.-dec. 2011), 47–59.

[16] Hunter, J. D. Matplotlib: A 2D graphics environment. *Computing in Science and Engineering 9*, 3 (May-Jun 2007), 90–95.

[17] Jones, E., Oliphant, T., Peterson, P., et al. SciPy: Open source scientific tools for Python, 2001–2013.

[18] Lange, M. On the Uncertainty of Wind Power Predictions—Analysis of the Forecast Accuracy and Statistical Distribution of Errors. *Journal of Solar Energy Engineering 127*, 2 (2005), 177–184.

[19] Pinson, P., and Girard, R. Evaluating the quality of scenarios of short-term wind power generation. *Applied Energy 96*, 0 (2012), 12–20.

[20] Pinson, P., and Madsen, H. Adaptive modelling and forecasting of offshore wind power fluctuations with Markov-switching autoregressive models. *Journal of Forecasting 31*, 4 (2012), 281–313.

[21] Pinson, P., Madsen, H., Nielsen, H. A., Papaefthymiou, G., and Klöckl, B. From probabilistic forecasts to statistical scenarios of short-term wind power production. *Wind Energy 12*, 1 (2009), 51–62.

[22] Pinson, P., Papaefthymiou, G., Klöckl, B., and Verboomen, J. Dynamic sizing of energy storage for hedging wind power forecast uncertainty. In *Power Energy Society General Meeting, 2009. PES '09. IEEE* (july 2009), pp. 1–8.

[23] R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2009.

[24] Söder, L. Integration study of small amounts of wind power in the power system. Tech. rep., 1994.

[25] Söder, L. Simulation of wind speed forecast errors for operation planning of multiarea power systems. In *International Conference on Probabilistic Methods Applied to Power Systems, 2004* (Sept. 2004), pp. 723–728.

[26] Teleke, S., Baran, M., Huang, A., Bhattacharya, S., and Anderson, L. Control Strategies for Battery Energy Storage for Wind Farm Dispatching. *IEEE Trans. Energy Convers. 24*, 3 (sept. 2009), 725–732.

[27] Tsay, R. S. Model Checking via Parametric Bootstraps in Time Series Analysis. *Journal of the Royal Statistical Society. Series C (Applied Statistics) 41*, 1 (1992), 1–15.

[28] Ummels, B., Gibescu, M., Pelgrum, E., Kling, W., and Brand, A. Impacts of Wind Power on Thermal Generation Unit Commitment and Dispatch. *IEEE Trans. Energy Convers. 22*, 1 (Mar. 2007), 44–51.

[29] Yao, D., Choi, S., Tseng, K., and Lie, T. Determination of Short-Term Power Dispatch Schedule for a Wind Farm Incorporated With Dual-Battery Energy Storage Scheme. *IEEE Trans. Sustain. Energy 3*, 1 (Jan. 2012), 74–84.