

Energy Storage Control with Aging Limitation

Pierre Haessig^{*}, Hamid Ben Ahmed, Bernard Multon[†]

Abstract

Energy Storage Systems (ESS) are often proposed to mitigate the fluctuations of renewable power sources like wind turbines. In such a context, the main objective for the ESS control (its energy management) is to maximize the performance of the mitigation scheme.

However, most ESS, and electrochemical batteries in particular, can only perform a limited number of charge/discharge cycles over their lifetime. This limitation is rarely taken into account in the optimization of the energy management, because of the lack of an appropriate formalization of cycling aging.

We present a method to explicitly embed a limitation of cycling aging, as a constraint, in the control optimization. We model cycling aging with the usual “exchanged energy” counting method. We demonstrate the effectiveness of our aging-constrained energy management using a publicly available wind power time series. Day-ahead forecast error is minimized while keeping storage cycling just under an acceptable target level.

1 Introduction to Aging Control

1.1 Why Limiting Aging in Energy Storage Control?

Energy Storage Systems (ESS) like electrochemical batteries can be used as buffers to reduce power fluctuations in different applications. For wind or solar power generation, an ESS can mitigate the fluctuations of the renewable power source. In an hybrid electric vehicle (HEV), it absorbs the fast fluctuations of the driving power profile to keep the combustion engine close to an efficient operating point.

^{*}P. Haessig is with IETR, CentraleSupélec, Rennes, France. This work was partly conducted while he was at ENS Rennes, sponsored by EDF R&D.

[†]H. Ben Ahmed, B. Multon are with SATIE lab, ENS Rennes, Bruz, France

In these two applications, the control of the ESS is often treated as an optimization problem where the objective is to maximize the performance of the overall system. Performance criterions can be, for example, the fuel efficiency for a HEV, or a measure of the fluctuations at the output of a combined wind-storage system (cf. figure 1 introduced in the next section). The control decides how and when to charge and discharge the battery to maximize the performance objective.

However, most ESS, and electrochemical batteries in particular, can only perform a limited number of charge/discharge cycles over their lifetime. This aging phenomenon causes a degradation of the battery parameters: decrease of the capacity, increase of the series resistance. This can lead to an eventual *ESS replacement*. This is why aging is crucial for evaluating and then minimizing the *life-cycle cost* of an ESS [1,2].

Unfortunately, battery aging is seldom taken *directly* in the control optimization. Often, it is only *after* simulating the behavior of the system, that an aging study is conducted. Before discussing existing work on aging control, we need to give an overview of aging modeling.

1.2 Models for Cycling Aging

The purpose of limiting battery aging requires a model for the aging phenomenon. Battery aging can be studied at the microscopic scale of the degradation processes, which is a research field by itself. However, for control, we need simpler models that are defined on the system level. Our work relies on aging models from the literature which are well in line with the datasheets provided by battery manufacturers. The verification of such models using (accelerated) aging test benches is again another field of expertise.

A commonly used empirical model for cycling ag-

ing is the “Ah throughput” model. It is widely described in the literature [3, §4] [4, §4], with some variations like “weighted Ah throughput” models to account for technology-specific aggravating aging factors.

This model considers that a battery can exchange a fixed amount of charge over its life. As such, the model consists in integrating over time the current that goes *in and out* of the battery (thus the name “Ah-throughput”). We use a common variation which consists in counting the exchanged energy instead of the charge (Wh vs. Ah). This *energy counting* model integrates the absolute power P_{sto} of the battery:

$$E_{exch}(t) = \int_0^t |P_{sto}| dt \quad (1)$$

and this exchanged energy is then compared with the energy exchanged during one full charge-discharge cycle, i.e. two times the energetic capacity of the battery ($2E_{rated}$). The ratio gives N_{cycl} , the number of *equivalent full cycles*:

$$N_{cycl}(t) = E_{exch}(t)/(2E_{rated}) \quad (2)$$

This number increases with time, and when it reaches N_{life} , the maximum number of full charge-discharge cycles, the battery is considered to be in end-of-life and should be replaced¹. N_{life} is typically between 500 and 5000 for batteries. An important property of this counting method is that it allows a number of cycles that is *inversely proportional* to the amplitude of these cycles (often called Depth of Discharge). This fact is often verified on the “aging curves” provided by manufacturers of lead-acid batteries. The reader can find a deeper analysis on the link between energy counting and cycle counting in Serrao *et al.* [5] and our thesis [6, §2.2].

Now that an aging model is established, we can turn to our main contribution: the mitigation of cycling aging through the ESS control.

1.3 Previous Work on Aging Control

The problem we have described in section 1.1 was already underlined by some authors [7–9]. We can summarize some keys aspects of their contributions: Serrao *et al.* [7] (HEV context) and Borhan *et al.*

[8] (wind power context) both use weighted charge counting to model cycling aging, while Koller *et al.* [9] (peak shaving context) use a piece-wise affine model with quadratic cost because it enables efficient optimization solving.

All these approaches attempt to reduce the aging by adding the aging increase *as a penalty* in the optimization criterion, which otherwise contains only a performance criterion (like the HEV fuel efficiency). The control designer is thus forced to *tune a weighting factor* to find a satisfactory compromise between performance and aging.

We propose instead to express aging limitation as an *inequality constraint* rather than a cost penalty. This can ease the implementation of aging limitation, because the control designer can directly set a desired maximum number of cycles N_{life} , with no need for tuning.

1.4 Reformulating aging limitation as a constraint

Based on the cycle counting model (2), we can indeed express aging limitation as an inequality constraint. Given T_{life} , the expected lifetime of the project involving the ESS, the non replacement of the storage during the operational period is expressed by $N_{cycl}(T_{life}) \leq N_{life}$. This translates into a constraint on the *mean absolute storage power*:

$$\frac{1}{T_{life}} \int_0^{T_{life}} |P_{sto}| dt = \langle |P_{sto}| \rangle_{T_{life}} \leq P_{exch} \quad (3)$$

where P_{exch} is what we call the *mean exchangeable power*. It is the ratio of the lifetime exchangeable energy of the storage with the duration of the project:

$$P_{exch} = 2E_{rated}N_{cycl}/T_{life} \quad (4)$$

Inequality (3) is the condition that should be satisfied by the ESS control algorithm (which sets P_{sto} at each time) to limit the aging. However, this is an integral constraint on a very long time horizon (T_{life} typically in the range of 5 to 20 years). This horizon is much longer than a usual ESS control horizon (on the order of the energy/power ratio of the ESS, i.e. minutes to hours). It cannot be practically solved by common control algorithms such as Model Predictive Control (MPC) or Stochastic Dynamic Programming (SDP). As such, expressing the lifetime constraint (3)

¹A common and equivalent expression is to define a State of Aging (SoA) as the $N_{cycl}(t)/N_{life}$ ratio. SoA starts at 0 zero, and end-of-life is reached when $SoA(t) = 1$

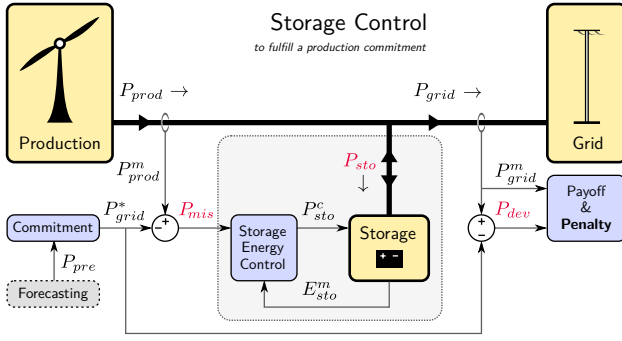


Figure 1: Wind-storage system used as a context for storage aging limitation

is only a first step, and our key contribution is to replace lifetime constraint (3) with a manageable constraint, in the next section.

2 Problem Description

2.1 Modeling

2.1.1 Wind-storage system

We base our study on the context of a wind-storage system for a day-ahead production commitment. The storage is used to mitigate the fluctuations of the wind power plant (see [10, 11] for previous work on this context). Figure 1 shows the variable and energy flows of this system, with the three main ones highlighted in red:

- P_{mis} is the difference between the wind power production P_{prod} and the commitment P_{grid}^* made one day in advance, based on a production forecast².
- P_{sto} is the power absorbed by the storage (convention $P_{sto} > 0$ when charging, $P_{sto} < 0$ when discharging).
- P_{dev} is the commitment deviation ($P_{grid} - P_{grid}^*$) for which the wind operator must pay penalties to the grid operator (cf. [10] for detailed context).

The commitment deviation can be written as:

$$P_{dev} = P_{mis} - P_{sto} \quad (5)$$

²see ANEMOS.plus project report [12] for a state of the art on wind power forecasting




Figure 2: Graphical representation of the dynamical models for the Energy Storage System and its aging. On the left, the usual stock of stored energy (6). On the right, the auxiliary stock of “exchangeable energy” (8), which we introduce to enforce aging limitation (3).

to highlight how the storage can be seen to act directly as a mitigation of the day-ahead forecast error P_{mis} .

2.1.2 Energy Storage System

For the purpose of energy management, we need a simple energetic model of the ESS. We use a discrete time model, with Δ_t as the time step:

$$E_{sto}(k+1) = E_{sto}(k) + (P_{sto}(k) - P_{losses})\Delta_t \quad (6)$$

where E_{sto} is the energy stored in the ESS. P_{losses} represents all the energy losses of the storage (in particular: self-discharge and Joule losses) and, in general, is a complex function of P_{sto} and E_{sto} , and depends on the technology. Since we do not focus on these losses here, we consider a lossless storage: $P_{losses} = 0$.

The amount of stored energy E_{sto} is constrained by the rated energy E_{rated} :

$$0 \leq E_{sto} \leq E_{rated} \quad (7)$$

and we can define the State of Energy: $SoE = E_{sto}/E_{rated} (\in [0,1])$. The dynamical model defined by (6) is graphically represented on the left side of figure 2.

2.1.3 Exchangeable Energy

As explained in section 1.4, aging limitation can be expressed as an integral constraint (3) on an horizon that is too long. We thus introduce a new auxiliary

state variable to embed constraint (3) in a manageable way. We call it the *exchangeable energy* stock, with the following dynamical behavior:

$$X_{sto}(k+1) = \text{sat}\{X_{sto}(k) + (P_{exch} - |P_{sto}(k)|)\Delta_t\} \quad (8)$$

where P_{exch} is the mean exchangeable power (4) and “sat” enables the *overflow* of this stock beyond a threshold X_{max} :

$$\text{sat}(x) = \begin{cases} x & \text{if } x \leq X_{max} \\ X_{max} & \text{if } x > X_{max} \end{cases} \quad (9)$$

We illustrate this dynamics on the right side of figure 2. A key stage is requiring this stock to be non empty: $0 \leq X_{sto}$, which translates into a constraint on the control variable P_{sto} :

$$|P_{sto}(k)| \leq P_{exch} + X_{sto}(k)/\Delta_t \quad (10)$$

This constraint is indeed simpler than (3) because it only involves the current value of the state variable X_{sto} . One can show that combining constraint (10) with the dynamical equation (8) gives a *sufficient condition* to respect the original aging limitation (3).

X_{sto} interpretation: We propose to interpret this new variable X_{sto} as a *buffer of exchangeable energy*. When the battery is heavily used ($|P_{sto}| \geq P_{exch}$), this stock decreases. When the battery is less used ($|P_{sto}| < P_{exch}$), the stock regenerates. If X_{sto} reaches zero, this forces $|P_{sto}(k)| \leq P_{exch}$, which is a conservative way to respect condition (3). The finite range of X_{sto} explains that in the long run, the average of $|P_{sto}|$ is indeed smaller than P_{exch} .

One downside of this formulation is that it introduces the parameter X_{max} , the size of the exchangeable energy stock, which must be chosen. If it is too small, constraint (10) falls back to $|P_{sto}(k)| \leq P_{exch}$ which is stricter than (3). A too big value yields a big stock that the control optimization must manage, so it brings back the problem of a too long optimization horizon. At the end of the article (section 3.4), we extend this qualitative reasoning with numerical results on the effect choosing X_{max} . We argue that a “big enough” size gives enough freedom so that the system performance is eventually the same as with the original constraint (3).

2.1.4 Forecast Error Persistence

Day-ahead forecast error P_{mis} is a *stochastic input* for the ESS control. Also important, it exhibits *persistence* (e.g. positive correlation) along several hours

[10]. We capture both the randomness and the persistence using an Autoregressive AR(1) model:

$$P_{mis}(k+1) = \phi P_{mis}(k) + w(k) \quad (11)$$

where $w(k)$ is a Gaussian white noise.

This AR(1) model has two parameters: the standard deviation σ_P of P_{mis} (i.e. the RMS forecast error) which is linked to the variance of w , and the AR coefficient ϕ which is the correlation between two successive hours. Both must be estimated on actual time series [10].

2.2 Control Optimization

Now that we have expressed the system dynamics, we can formulate the control optimization problem. Our control objective is to keep the total output of the wind-storage system ($P_{grid} = P_{prod} - P_{sto}$) around the day-ahead commitment P_{grid}^* , in a $\pm P_{tol}$ interval. The control, acting on the storage power P_{sto} , should minimize *on average* a penalty at each instant:

$$J = \frac{1}{K} \mathbb{E} \left\{ \sum_{k=0}^{K-1} \text{cost}(k) \right\} \quad \text{with } K \rightarrow \infty \quad (12)$$

where $\text{cost}(k)$ is the following penalty function:

$$\text{cost}(k) = \max \{0, |P_{dev}(k)| - P_{tol}\} \quad (13)$$

which penalizes in absolute value each deviation outside the tolerance band. This penalty with a free tolerance band is inspired by a grid code for wind-storage systems in French islands [10] and can be similarly found in the Hungarian grid code [13, §5.A].

This optimization includes the temporal constraints introduced with the dynamical equations (6), (8) and (11). The expectation $\mathbb{E}\{\}$ is needed because of the stochastic input w in (11). Therefore, minimizing J is a *stochastic dynamic optimization* problem [14], which we solve using our open source Stochastic Dynamic Programming (SDP) package *stodynprog* [15]. Note that using SDP is *not a requirement* for solving the formulation of aging limitation we present here. Another control framework, like the popular Model Predictive Control (MPC, with Koller *et. al* work [9] as one example) could be used as well.




Figure 3: Wind-storage system fulfilling a day-ahead production commitment. Simulation with control C1: optimal ESS control with no aging limitation




Figure 4: Wind-storage system fulfilling a day-ahead production commitment. Simulation with control C3: optimal aging-constrained ESS control

3 Aging Control Results

3.1 Input Data

We use the publicly available NREL “Eastern Wind Dataset” [16] as a test case for our aging control method. It provides 3 years of production and day-ahead forecast data³ at an hourly timestep ($\Delta_t = 1$ h) for many wind farms in the US. We choose farm #7277 and normalize the powers by its production capacity (132.3 MW) so that production and forecasts are expressed in per unit (pu). This farm has a mean production of 0.343 pu and RMS forecast error is $\sigma_P = 0.195$ pu. The AR coefficient ϕ , which gives the correlation of forecast errors between two successive hours, is estimated at 0.79. This high positive temporal correlation is typically observed with day-ahead forecasts and it adversely impacts the system performance [10]. On figures 3 and 4, we represent a 10 days extract of this input data (forecast in gray, production in light blue).

3.2 Test Case

We consider a storage of capacity $E_{rated} = 1$ h⁴. Deviation tolerance P_{tol} is set to 0.2 pu (shaded area on figures 3 and 4). This 20% tolerance is in line with the

³this dataset is in fact *synthetic*: production and forecasts are reconstructed from numerical weather models [17].

⁴ E_{rated} is expressed in hours when working in per units. Actual capacity is 1 h × the rated power of the wind farm, which is a typical value in such a context of day-ahead production hedging.

grid code for French islands. For aging limitation, we require a maximum of $N_{life} = 3000$ equivalent full cycles over the period $T_{life} = 20$ years. This gives an exchangeable power $P_{exch} = 0.034$ pu, which is quite small compared to the RMS forecast error σ_P . We choose this (moderately) ambitious aging limitation so that it cannot be reached “by chance”, without an explicit action in the energy management. This way we can show the effectiveness of our aging control. Finally, we set X_{max} to 1.71 h ($P_{exch} 50$ h) but we explain this choice further, in section 3.4.

We simulate the wind-storage system with different controls and collect three performance statistics (reported in table 1):

- Aging: the number of equivalent full cycles $N_{cycle}(T_{life})$, which we would like to be less than $N_{life} = 3000$.
- Ovtol, shorthand for “Over tolerance”: the proportion of hourly time steps spent above the tolerance threshold ($|P_{dev}(k)| > P_{tol}$).
- Ovtol MAE: Mean Absolute Error above the tolerance threshold. This is in fact the criterion J minimized by the control, given penalty function (13).

3.3 Simulation Results

For each control strategy, we run the simulation with the entire 3 years dataset and we report the statistics

in table 1. The controls we compare, along their respective results, are:

1. C1: optimal control with no aging limitation. This gives the best performance for J (0.013 pu), but aging is twice above the target of 3000 cycles/20 yr.
2. C2: we overload C1 with our aging limitation (10) introducing X_{sto} in the dynamics. Aging limitation is indeed effective, but the performance is severely decreased because C1 doesn't anticipate X_{sto} evolution. As a consequence, X_{sto} value is often zero, or close to, which implies a overly conservative limitation of $|P_{sto}|$ through inequality (10).
3. C3: optimal control with aging limitation (using X_{sto}): Aging limit is respected, and the performance (0.014 pu) is not far from the best performance without aging limitation (C1).

As a baseline reference, C0 gives the performance statistics in the absence of storage. These simulations show the effectiveness of our aging control C3 at keeping cycling aging just below the user imposed limit, while still keeping the best possible performance, like the regular optimal control C1.

Table 1: Performance comparison

Control	Aging and performance statistics		
	Aging	Ovtol	Ovtol MAE
C0	–	27.5 %	0.032 pu
C1	6 372	8.49 %	0.013 pu
C2	2 998	19.97 %	0.023 pu
C3	2 966	9.97 %	0.014 pu

To illustrate the qualitative behavior of these two controls, we represent an extract of trajectories with C1 and C3 on figures 3 and 4 respectively. The time periods when $|P_{dev}| > P_{tol}$ are highlighted in orange. One can see that the time spent outside the tolerance band is quite similar between the two. Looking at the SoE (bottom panels), there are more fluctuations for C1 than C3, but the slow (daily) variations are similar. However, we underline that strategy C3 cannot be reduced to a simple linear low-pass filtering of C1. We can observe that with C3, the battery spends more time at rest ($P_{sto} = 0$ so $P_{grid} = P_{prod}$ and $SoE(t)$ is

constant). All in all, the “thrifter” battery management C3 is responsible for the reduction in the number of cycles by a factor of two compared with C1.

3.4 Choosing the Aging Control Horizon

We have so far left undiscussed the choice of the only tunable parameter of our method: X_{max} , the size of the exchangeable energy stock X_{sto} . Let us remind that this stock is used to give freedom to the control to allocate more battery power (i.e. consume cycles) when needed, while still keeping $|P_{sto}|$ below P_{exch} on average.

For better reasoning, we express this stock size as a time T_X using the relation

$$X_{max} = P_{exch} \times T_X \quad (14)$$

Parameter T_X is the time it takes to recharge the entire stock, starting from zero (X_{sto} stock is indeed recharged, according to (8), at a rate equal to P_{exch}). Thus, from the perspective of aging control, it represents the time horizon on which the ESS control algorithm can borrow exchangeable energy from the future. Therefore, we call T_X the *aging control horizon*.

We study its effect by varying its value from zero to 500 hours in our age limiting control C3. We show the resulting performance statistics on figure 5 (red curve). For a too small horizon ($T_X < 1$ h), there is a major performance loss due to the conservativeness of our aging limitation method. As the horizon is increased, the performance loss decreases. Then, for an horizon T_X greater than about 100 hours, there is an optimal plateau as $T_X \rightarrow \infty$. We thus claim that our method becomes equivalent to an optimal control satisfying the initial aging constraint (3) that spans over T_{life} (20 years here).

These numerical results support the qualitative reasoning proposed at the end of section 2.1.3. Also, this justifies a posteriori our choice of $T_X = 50$ hours in our previous simulations: it yields an (almost) optimal performance while being small enough to ensure the fastest convergence of the SDP optimization algorithm⁵.

Finally, we can observe on figure 5 a performance gap of about 0.001 pu between C3 and C1 (control

⁵more precisely, we use a “policy iteration” algorithm [14, 15]. It includes iterations along time in each “policy evaluation” step, and the number of those iterations is dictated by the time constants of all stocks, e.g. T_X .




Figure 5: Effect of the aging control horizon on the performance of the optimal aging-constrained control C3. A too small T_X (under ~ 100 h) yields a too conservative aging limitation which decreases the performance. As a baseline, control C1 with no aging limitation shows the performance difference due to the aging limitation constraint.

without aging limitation, blue line). We claim that this is the *price of the aging constraint* (a constrained minimum is always worse or equal than a constraint-free one). A further study could be to vary the aging limit (3000 cycles on 20 years in this article) to generate a trade-off curve between the performance and the limit. This would outline a Pareto front between these two conflicting objectives (minimizing output deviations and cycling aging).

4 Conclusion

We introduced a formulation of cycling aging (based on exchanged energy counting) with two advantages: it fits naturally in the ESS control optimization and it enables the control designer to directly set a maximum number of battery cycles over the project lifetime. We illustrated the effectiveness of our scheme with a simulation on an open dataset. On this example, cycling is reduced by 50 % with a less than 10 % decrease in performance.

Further studies could include looking at the *trade-off* between the limitation of storage aging and the performance. Also, we plan to adapt our formalism to other aging models, in particular weighted charge counting [3,4], and more importantly *calendar aging*. Calendar aging can be included in control optimization when it depends on operational conditions like

SoE (particularly true for super capacitors [1]).

Acknowledgment

We are grateful to NREL for providing the Eastern Wind dataset [16,17] which is very useful as an open benchmark for simulations of wind-storage systems.

References

- [1] T. Kovaltchouk, B. Multon, H. Ben Ahmed, J. Aubry, and P. Venet, “Enhanced Aging Model for Supercapacitors taking into account Power Cycling: Application to the Sizing of an Energy Storage System in a Direct Wave Energy Converter,” in *9th Ecological Vehicles and Renewable Energies (EVER) 2014 conference, Monte Carlo, Monaco*, Mar. 2014.
- [2] J. Aubry, P. Bydlowski, B. Multon, H. Ben Ahmed, and B. Borgarino, “Energy Storage System Sizing for Smoothing Power Generation of Direct Wave Energy Converters,” in *3rd International Conference on Ocean Energy*, 2010.
- [3] H. Bindner, T. Cronin, P. Lundsager, J. F. Manwell, U. Abdulwahid, and I. Baring-Gould, “Lifetime Modelling of Lead Acid Batteries,” Risø National Laboratory, Tech. Rep., Apr. 2005.
- [4] D. U. Sauer and H. Wenzl, “Comparison of different approaches for lifetime prediction of electrochemical systems—Using lead-acid batteries as example,” *Journal of Power Sources*, vol. 176, no. 2, pp. 534–546, 2008.
- [5] L. Serrao, S. Onori, G. Rizzoni, and Y. Guezennec, “A novel model-based algorithm for battery prognosis,” in *Proceedings of the Seventh IFAC Symposium on Fault Detection, Supervision and Safety of Technical Processes (SAFEPROCESS 09)*, 2009.
- [6] P. Haessig, “Dimensionnement & gestion d’un stockage d’énergie pour l’atténuation des incertitudes de production éolienne,” Ph.D. dissertation, ENS Cachan, Jul. 2014.
- [7] L. Serrao, S. Onori, A. Sciarretta, Y. Guezennec, and G. Rizzoni, “Optimal energy management of hybrid electric vehicles including

- battery aging,” in *American Control Conference (ACC), 2011*, June 2011, pp. 2125–2130.
- [8] H. Borhan, M. A. Rotea, and D. Viassolo, “Optimization-based power management of a wind farm with battery storage,” *Wind Energy*, vol. 16, no. 8, pp. 1197–1211, Nov. 2013.
- [9] M. Koller, T. Borsche, A. Ulbig, and G. Andersson, “Defining a degradation cost function for optimal control of a battery energy storage system,” in *IEEE PowerTech 2013 Conference, Grenoble, France*, Jun. 2013.
- [10] P. Haessig, B. Multon, H. Ben Ahmed, S. Lascaud, and P. Bondon, “Energy storage sizing for wind power: impact of the autocorrelation of day-ahead forecast errors,” *Wind Energy*, vol. 18, no. 1, pp. 43–57, Jan. 2015, published online Oct 2013.
- [11] P. Haessig, B. Multon, H. Ben Ahmed, S. Lascaud, and L. Jamy, “Aging-aware NaS battery model in a stochastic wind-storage simulation framework,” in *IEEE PowerTech 2013 Conference, Grenoble, France*, Jun. 2013.
- [12] G. Giebel, R. Brownsword, G. N. Kariniotakis, M. Denhard, and C. Draxl, “The state-of-the-art in short-term prediction of wind power: A literature overview,” ANEMOS.plus, Tech. Rep., 2011.
- [13] B. Hartmann and A. Dán, “Cooperation of a Grid-Connected Wind Farm and an Energy Storage Unit – Demonstration of a Simulation Tool,” *IEEE Trans. Sustain. Energy*, vol. 3, no. 1, pp. 49–56, Jan. 2012.
- [14] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 3rd ed. Athena Scientific, 2005.
- [15] P. Haessig, T. Kovaltchouk, B. Multon, H. Ben Ahmed, and S. Lascaud, “Computing an Optimal Control Policy for an Energy Storage,” in *6th European Conference on Python in Science (EuroSciPy 2013), Brussels, Belgium*, Aug. 2013, pp. 51–58.
- [16] National Renewable Energy Laboratory, “Eastern Wind Integration and Transmission Study,” Tech. Rep. NREL/SR-5500-47078, Feb. 2011.
- [17] AWS Truewind, “Development of Eastern Regional Wind Resource and Wind Plant Output Datasets: March 3, 2008 - March 31, 2010,” Tech. Rep. NREL/SR-550-46764, Dec. 2009.