

Using Approximate Dynamic Programming for Estimating the Revenues of a Hydrogen-based High-Capacity Storage Device

Vincent François-Lavet*, Raphael Fonteneau*, Damien Ernst*

*Department of Electrical Engineering and Computer Science
University of Liège, Belgium

Email: {v.francois, raphael.fonteneau, dernst}@ulg.ac.be

Abstract—This paper proposes a methodology to estimate the maximum revenue that can be generated by a company that operates a high-capacity storage device to buy or sell electricity on the day-ahead electricity market. The methodology exploits the Dynamic Programming (DP) principle and is specified for hydrogen-based storage devices that use electrolysis to produce hydrogen and fuel cells to generate electricity from hydrogen. Experimental results are generated using historical data of energy prices on the Belgian market. They show how the storage capacity and other parameters of the storage device influence the optimal revenue. The main conclusion drawn from the experiments is that it may be advisable to invest in large storage tanks to exploit the inter-seasonal price fluctuations of electricity.

I. INTRODUCTION

Developing sustainable energy systems is one of the most critical issues that today's society must address. Due to the fluctuating nature of the Renewable Energy Sources (RES) generation, fossil-fuel-based generation capacity is currently still needed to provide flexibility and to adequately cover peak demand. This issue can be partially addressed or mitigated in several ways [1]: (i) by diversifying the types of renewable energy sources to reduce the correlation between the amount of energy supplied by these sources, which lowers the risk of shortage of supply, (ii) by developing electricity storage capacity, (iii) by increasing the flexibility of the demand, to smooth out peak demand or (iv) by developing the electrical network since the variance in the energy supplied by renewable sources tends to decrease with the size of the zone on which they are collected [2]. This has been the main motivation for developing the European network in recent years. Note that authors have also reported that the variance in the energy supplied by renewables could be further decreased by building a global electrical grid that connects continents together [3], [4].

During recent years, storage has gradually become more and more profitable thanks to technological progress. Consequently, economic actors on the energy market are currently planning to invest additional funds in storage devices. Among the different storage technologies, pumped-storage hydroelectricity and batteries are currently among the most mature. Other technologies exist such as for example supercapacities, energy conversion to natural gas, compressed air energy storage, flywheels, superconducting magnetic energy storage and storage of electricity in the form of hydrogen. This latter one seems to be particularly promising due to its

capability to store large quantities of energy at a relatively low cost, and is therefore well suited for long-term storage [5]. Additionally, the round trip efficiency of hydrogen-based storage devices is rather good. For example, the energy efficiency of an electrolyzer is around 80% and the one of a fuel cell is generally between 40% and 60%, which results in an overall round-trip efficiency of 35% up to 50%, with the potential to get an efficiency higher than 70% in hybrid fuel cell/turbine systems and more than 80% in Combined Heat and Power (CHP) systems.

However, before investing in such a hydrogen-based storage technology, a careful analysis of the return on investment needs to be carried out. Such an analysis implies, among others, to be able to estimate the revenues that can be generated by such a storage device on the power exchange markets, which is the focus of this paper. We will consider the case of a company that operates the hydrogen-based high-capacity storage device and makes money by buying or selling electricity on the day-ahead market. In such a context, the company has to decide on the day-ahead which amount of electricity to store or to generate for every market period. The main complexity of this decision problem originates from the fact that a decision to store or generate electricity at one specific market period may not only significantly impact the revenues that could be generated at other market periods of the day, but also the revenues that could be generated months ahead. As a result, long optimization horizons have to be considered for computing operation strategies for high-capacity storage devices.

The valuation of energy storage technologies on power markets has already received considerable attention in the scientific literature [6], [7], [8], [9], [10], [11]. For example, reference [10] proposes an approach based on mixed-integer programming for optimizing bidding strategies for hydropower. This approach can handle uncertainty in market prices and water inflows. However, the computational complexity of this technique grows very rapidly with the state/action space, which makes this approach unsuitable for estimating the revenues that can be generated by a storage capacity over an extended period of time. Another example is reference [11] where a methodology based on Approximate Dynamic Programming (ADP) is proposed for jointly optimizing in the day-ahead the trading of renewable energy and of the storage management strategies.

Before explaining the details of this approach, we will describe in Section II the bid process for a typical day-ahead electricity market such as the Belgian electricity market and lay out in Section III a mathematical model for energy storage utilizing hydrogen. A first formulation of our problem as a dynamic programming problem will be stated in Section IV where we assume that only the market prices of the next market day are known. Section V specifies this formulation to the case where the market prices are assumed to be known over the whole optimization horizon and provides a fully specified algorithm that exploits this new formulation for computing the maximum operational revenue. The computational complexity of this algorithm is linear with respect to the cardinality of the discretized state space, the cardinality of the discretized action space, and the optimization horizon. Section VI provides experimental results computed from historical data gathered over the Belgian electricity market. Finally, Section VIII concludes the paper.

II. OPTIMIZATION ON THE DAY-AHEAD ENERGY MARKET

Let us consider a power exchange market for the day-ahead trading of electricity, providing the market with a transparent reference price. Producers and retailers submit each day offers to the market operator for the day-ahead. An offer is defined by a volume and a limit price, and can span several market periods. The market clearing price is computed by the market operator at the intersection of the supply and the demand curves. The prices for electricity on the Belgian day-ahead market are determined via a blind auction with the possibility to define linked Block Orders that allow the execution of profile blocks to be subjected to the execution of other blocks. This possibility allows for the design of complex linked structures (i.e. families) that take into account the different possible price outcomes of the market clearing price. Figure 1 shows the distribution of prices over the year 2013.

In this paper, we consider that the storage capacity is an agent which interacts with the electricity exchange market under the following assumptions:

- the evolution of the price of electricity does not depend on the behavior of this agent. This hypothesis is equivalent to assuming that the bids for supply or demand from the actor do not change the market clearing price significantly¹.
- the evolution of the prices is known when determining the agent behavior.

III. PROBLEM FORMALIZATION

Let us introduce a discrete-time system whose state variable is fully described by the amount of energy in the storage device. The state space \mathcal{S} contains all possible states $s_{i,j} \in \mathcal{S}$, where the indices (i, j) refer to hour j during day i (in MWh). Let \mathcal{A} be the set of possible actions and $\vec{a}_{i,j} \in \mathcal{A}$ the action taken at time (i, j) . At every time step, an action $\vec{a}_{i,j} = [a_{i,j}^{GR}, a_{i,j}^{RG}] \in \mathcal{A}$ is applied on the system, where $a_{i,j}^{GR}$ is the amount of energy transferred into the storage (R) from

¹That effect is measured by the market resilience which is the price sensitivity due to an increase in offer or demand on the market. It is of the order of $5 \cdot 10^{-3} \text{€}/\text{MWh}$ on the Belgian power exchange market [12].




Fig. 1. Histogram of the price of electricity for the year 2013.

the grid (G), and $a_{i,j}^{RG}$ is the amount of energy taken out of the storage (R) to the grid (G). The actions $a_{i,j}^{GR}$ and $a_{i,j}^{RG}$ are non-negative. The considered dynamics is defined over n_D days and n_H market periods ($n_H = 24$). We denote by I and J the sets of time indices:

$$\begin{aligned} I &= \{0, \dots, n_D - 1\}, \\ J &= \{0, \dots, n_H - 1\}. \end{aligned}$$

The system dynamics is given by the following equation:

$$\forall i \in I, \forall j \in J, \quad s_{i,j+1} = f(s_{i,j}, \vec{a}_{i,j}) \quad (1)$$

where we use the convention $s_{i,n_H} = s_{i+1,0}$ for any $i \in I$. The notation $t_{i,j}$ is introduced as the time index corresponding to time $(i, j) \in I \times J$. This transition function can be rewritten as follows:

$$s_{i,j} = s_{i,0} + \sum_{t=t_{i,0}}^{t_{i,j}-1} (a_t^{GR} - a_t^{RG}), \quad \forall (i, j) \in I \times J. \quad (2)$$

At any time $(i, j) \in I \times J$, the following constraints have to be satisfied:

$$s_{i,0} + \sum_{t=t_{i,0}}^{t_{i,j}-1} (a_t^{GR} - a_t^{RG}) \leq R^c \quad (3)$$

$$s_{i,0} + \sum_{t=t_{i,0}}^{t_{i,j}-1} (a_t^{GR} - a_t^{RG}) \geq 0 \quad (4)$$

where R^c is the energy capacity of the device.

The bidding process occurs only once for each day $i \in I$, which means that all actions taken on day $i+1$ are computed on day i . We denote by $\vec{s}_i \in \mathcal{S}$ the vector of states defined as $\vec{s}_i = [s_{i,0}, s_{i,1}, \dots, s_{i,n_H-1}]$. We denote by $A_i \in \mathcal{A}$ the matrix of actions defined as follows: $A_i = [\vec{a}_i^{GR}; \vec{a}_i^{RG}]$ with

$$\vec{a}_i^{GR} = [a_{i,0}^{GR}, a_{i,1}^{GR}, \dots, a_{i,n-1}^{GR}]$$

and

$$\vec{a}_i^{RG} = [a_{i,0}^{RG}, a_{i,1}^{RG}, \dots, a_{i,n-1}^{RG}].$$

The dynamics corresponding to the bidding process logic is then

$$s_{i+1,0} = F(s_{i,0}, A_i), \quad \forall i \in \mathbf{I}, \forall A_i \in \mathbf{A}_i \quad (5)$$

where the feasible action space \mathbf{A}_i is the set of matrices of actions A_i which satisfy the constraints at time $i \in \mathbf{I}$ defined by Equations (3) and (4).

We define a reward function $\rho(s_{i,0}, A_i, \vec{p}_i)$ for day i which measures the revenues generated by taking a sequence of actions A_i when starting from the state $s_{i,0}$, function of the vector of prices of electricity \vec{p}_i for day i . The value of the reward function is given by the total amount of money paid or collected when transferring energy to and from the grid. For every day i , the reward function is defined by

$$\rho(s_{i,0}, A_i, \vec{p}_i) = \sum_{t=t_{i,0}}^{t_{i,0}+n-1} r(\vec{a}_t, p_t)$$

where $r(\vec{a}_t, p_t)$ is given by

$$r(\vec{a}_t, p_t) = \left(a_t^{RG} \eta^d - \frac{a_t^{GR}}{\eta^c} \right) p_t$$

with η^d and η^c being the discharge and charge efficiencies, respectively.

In the context of the day-ahead energy market developed in Section II, the prices of electricity are known one day before, i.e. the prices of electricity on day $i \in \mathbf{I}$ are known when choosing the sequence of actions $A_i \in \mathbf{A}_i$. An admissible policy $\pi(i, s_{i,0}) : \mathbf{I} \times \mathcal{S} \rightarrow \mathbf{A}$ is a function that maps states into actions such that, for any state $s_{i,0}$, the action $\pi(i, s_{i,0})$ satisfies the constraints (3) and (4) (which defines the set of feasible actions $\mathbf{A}_i \subset \mathbf{A}$). We denote by Π such a set:

$$\Pi = \{ \pi : \mathbf{I} \times \mathcal{S} \rightarrow \mathbf{A} : \forall s_{i,0} \in \mathcal{S}, \forall i \in \mathbf{I}, \pi(i, s_{i,0}) \in \mathbf{A}_i \}$$

Arguably, the decision of a policy π to be made during the bidding process is whether to buy or sell energy to maximize the revenues on the long term. An optimal value function $V_{i+1}^*(s_{i+1,0})$ is introduced as the maximum expected revenue that can be obtained from time $(i+1, 0) = (i, n_H)$ over the remaining time-steps:

$$\forall s_{i+1,0} \in \mathcal{S}, V_{i+1}^*(s_{i+1,0}) = \max_{(A_{i+1}, \dots, A_{n_D-1}) \in \mathbf{A}_{i+1} \times \dots \times \mathbf{A}_{n_D-1}} \mathbb{E}_{\vec{p}_{i+1}, \dots, \vec{p}_{n_D-1}} \left[\sum_{k=i+1}^{n_D-1} \rho(s_{k,0}, A_k, \vec{p}_k) \right]$$

From these value functions, an optimal policy $\pi^* \in \Pi$ can be defined as follows:

$$\begin{aligned} & \forall i \in \mathbf{I}, \forall s_{i,0} \in \mathcal{S}, \\ & \pi^*(i, s_{i,0}) \in \arg \max_{A_i \in \mathbf{A}_i} \left(\rho(s_{i,0}, A_i, \vec{p}_i) + V_{i+1}^*(s_{i+1,0}) \right) \end{aligned} \quad (6)$$

IV. A DYNAMIC PROGRAMMING APPROACH TO COMPUTE THE OPTIMAL REVENUE OF STORAGE

In this paper, we make the (strong) assumption that the evolution of the prices is perfectly known. This has the two following consequences on the resolution of the above-described problem: (i) the problem becomes deterministic and (ii) the day-ahead structure of the problem disappears.

Let $Q_0, Q_1, \dots, Q_{24*n_D-1}$ be the sequence of functions defined as follows:

$$\begin{aligned} & \forall (s, a) \in \mathcal{S} \times \mathcal{A}, t = 0 \dots n_D * 24 - 1, \\ & Q_t(s, a) = r(s, a, p_t) + \max_{\text{feasible } a' \in \mathcal{A}} Q_{t+1}(f(s, a), a') \end{aligned} \quad (7)$$

with

$$Q_{n_D*24}(s, a) = 0, \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}.$$

It is straightforward to see that when the prices are known we have:

$$V_i^*(s) = \max_{a \in \mathbf{A}_i} Q_{i*24}(s, a)$$

From the sequence of functions Q_t , it is possible to estimate in a straightforward way the maximum revenue that can be generated by our storage capacity. We suggest to approximate the computation of this sequence of functions by discretizing the state and the action space [13], [14]. More specifically, the state space is discretized into a set $\{\sigma^{(i)}, i = 1 \dots n_S\}$, and the action space is discretized into a set $\{\alpha^{(i)}, i = 1 \dots n_A\}$. We also choose a projection function $\Gamma : \mathcal{S} \rightarrow \{\sigma^{(1)}, \dots, \sigma^{(n_S)}\}$ which projects any element of the state space \mathcal{S} into a unique element of the discretized space. In such a context, the problem is reduced to a dynamic programming problem with a finite horizon of $n_D * 24$ time-steps that can be solved with a backward value iteration algorithm [15]. The resulting algorithm is sketched in Procedure 1. It has a complexity proportional to the product of the size of the state space, the action space and the optimization horizon. $\mathfrak{A}(\sigma)$ denotes the set of feasible discretized actions for a given discretized state σ so that the maximization over possible actions $\alpha^{(i)}$ takes into account the constraints stated in Equations (3) and (4).

From the sequence of \widehat{Q}_t functions outputted by Procedure 1, one can extract a bidding policy. Note that the near-optimal revenue that is obtained from an initial state s_0 can be calculated as follows:

$$\arg \max_{\alpha' \in \mathfrak{A}(\Gamma(s_0))} \widehat{Q}_0(\Gamma(s_0), \alpha')$$

Another way to calculate this revenue is to simulate the system with the policy extracted from these \widehat{Q}_t functions. As way of example, Procedure 2 provides a way for computing the sequence of actions outputted by this policy when the initial state of the system is s_0 .

V. MATHEMATICAL MODEL FOR ENERGY STORAGE UNDER THE FORM OF HYDROGEN

Each storage capacity is defined by its maximum capacity, its maximum power consumption and restitution to the network

Algorithm 1 Q-iteration in the discretized state-action space

Require: $p_t, \forall t = 0, \dots, n_D * 24 - 1$;

```
for  $t = n_D * 24 - 1$  to 0 do {Backward loop over all time periods}
  for  $\sigma = \sigma^{(1)} \dots \sigma^{(n_S)}$  do {Loop over discretized states}
    for  $\alpha = \alpha^{(1)} \dots \alpha^{(n_A)}$  do {Loop over actions}
       $\hat{Q}_t(\sigma, \alpha) = r(\sigma, \alpha, p_t) + \max_{\alpha' \in \mathfrak{A}(\sigma')}$   $\hat{Q}_{t+1}(\sigma', \alpha')$  where
       $\sigma' = \Gamma(f(\sigma, \alpha))$ 
    end for
  end for
end for
return  $\hat{Q}_t, \forall t \in 0, \dots, n_D * 24 - 1$ 
```

Algorithm 2 Computation of the sequence of actions generated by the bidding policy

Require: $\hat{Q}_t, \forall t \in 0, 1, \dots, n_D * 24 - 1$; s_0

```
 $\sigma_0 = \Gamma(s_0)$ 
for  $t = 0$  to  $n_D * 24 - 1$  do {Loop over all time periods}
   $\alpha_t^* = \arg \max_{\alpha' \in \mathfrak{A}(\sigma_t)} \hat{Q}_t(\sigma', \alpha')$ 
   $\sigma_{t+1} = f(\sigma_t, \alpha_t^*)$ 
end for
return  $\alpha_t^*, \forall t \in 0, \dots, n_D * 24 - 1$ 
```

as well as the efficiencies for those three steps. A hydrogen-based high-capacity storage device is composed from three main parts: (i) an electrolyzer that transforms water into hydrogen using electricity (ii) a tank where the hydrogen is stored (iii) a fuel cell where the hydrogen is transformed into electricity. Figure 2 gives a schematic representation of such a device, whose main 3 elements are detailed hereafter.




Fig. 2. Sketch of the hydrogen based high-capacity storage.

A. Electrolysis

Currently the dominant technology for direct production of hydrogen (95%) is steam reforming from fossil fuels. However sustainable techniques also exist, such as electrolysis of water using electricity from one of the many renewable sources. It also has the advantage of producing high-purity hydrogen (>99.999%).

a): The technical performance of this process has a strong dependency on the rate at which the electrolysis is forced. The charge energy efficiency as a function of the cell


voltage is given by:

$$\eta^c = \frac{1.48}{\text{CellVoltage}}$$


The minimum voltage necessary for electrolysis is 1.23 V. Henceforth, the process can theoretically reach efficiencies above 100% but the rate at which the reaction happens is then very low [16]. The part of the voltage that exceeds 1.23 V is called overpotential or overvoltage, and leads to losses in the electrochemical process while allowing a higher rate in the reaction. Current density as a function of voltage is approximated at standard temperature for Flat-Plate Bifunctional Cells by

$$I = s \times (\text{CellVoltage} - 1.48),$$

where s is a constant dependent on the setup used for the electrolysis. The evolution of the efficiency with the voltage and with the power generated can be seen on Fig. 3(a) and 3(b), respectively.



(a) Efficiency as a function of Voltage



(b) Efficiency as a function of Power

Fig. 3. Evolution of the efficiency of the electrolysis process as a function of the rate at which the electrolysis is forced. Parameters used can be found in Table I.

B. Fuel cell

A fuel cell is a device that converts the chemical energy from a fuel, in this case hydrogen, into electricity through a chemical reaction with oxygen or another oxidizing agent. Unlike heat engine, the efficiency of a fuel cell is not limited

by the Carnot cycle and has a theoretical discharge efficiency $\eta^d = 83\%$ in the case of hydrogen. This efficiency is however lowered when the amount of power generated by the fuel cell increases as illustrated on Fig. 4. In standard operating conditions, the function $\eta^d(W_{fc})$ can be approximated as a linear equation:

$$\eta^d = \eta_{max}^d - s_{fc} W_{fc}$$

where s_{fc} is a constant dependent on the setup used for the fuel cell and W_{fc} is the power density of the fuel cell.

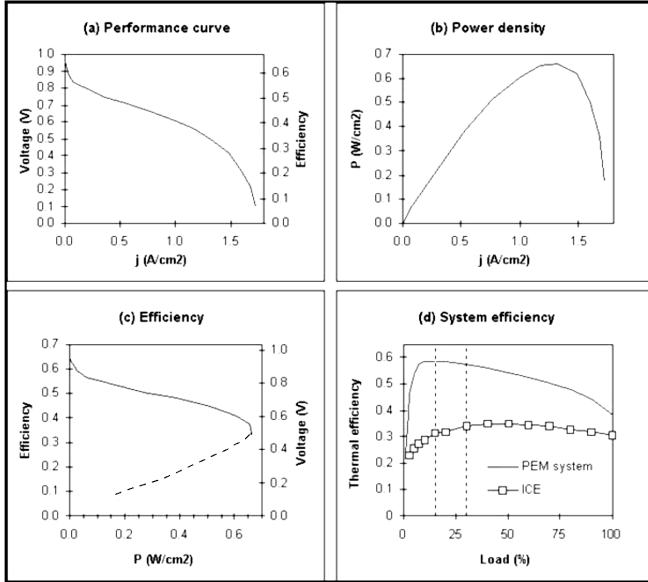


Fig. 4. Characteristics of a PEM fuel cell [17]. (a) The voltage as a function of the current density, commonly referred to as the performance curve. The efficiency is proportional to the voltage; it is indicated on the secondary vertical axis. (b) The power density as a function of the current density. (c) The efficiency as a function of the power density. The dotted line corresponds to the regime above maximum power. (d) The efficiency of a complete fuel cell system in a vehicle, as a function of power load, shown both for a PEM and an ICE. The vertical dotted lines indicate average loads in a car (left) and a bus or truck (right). The curves in (d) do not refer to the same fuel cell as in (a) to (c).

C. The storage device

One significant constraint that influences the choice of the storage device technology is often the energy density imposed by the application. In the case where hydrogen is to be used as a fuel stored on board of a vehicle, pure hydrogen gas must be pressurized or liquefied. The drawback is that it necessitates the use of external energy to power the compression. This constraint does not hold for grid energy storage, especially in the case where hydrogen can be stored in natural reservoirs such as in underground caverns, salt domes or depleted oil/gas fields.

In the following, the storage device will be characterized by the energy capacity of the device R^c (in MWh). It will be assumed that any leak in the storage device can be neglected.

VI. EXPERIMENTAL RESULTS

In the first part of this section, the algorithm described in Procedure 1 will be used to figure the maximum revenues

that could be generated over the period ranging from 2007 to 2013 by a high-capacity storage device whose parameters are defined in Table I. The historical data of electricity prices provided by Belpex over the last few years will be used as input [12]. In the second part, the influence of the discretization of the algorithm will be studied. Finally, the impact of the storage capacity on the overall gain will be analyzed.

A. Base case

To compute the revenues of the storage capacity defined by Table 1, we have first discretized the state-action space to be able to use Procedure 1. We choose for the state space a discretized step $\delta_s = 0.5$ MWh. The discretization step for the action space is taken equal to $\delta_u = 0.5$ MWh. That leads to a discretized state space equal to $\{0, 0.5, 1, \dots, R_c\}$ and a discretized action space equal to the finite set $\{-2, -1.5, -1, -0.5, 0, 0.5, 1, 1.5, 2\}$.

Electrolysis	$s_{electrolysis}$	1MA/V
Fuel cell	η_{max}^d	60%
	s_{fc}	0.4 MW^{-1}
	$W_{fc,max}$	5 MW
	$W_{fc,min}$	0.8 MW
Storage device	R^c	1000 MWh
	$s_{(0,0)}$	0 MWh

TABLE I. DATA USED FOR THE ELECTROLYSIS SETS AND FUEL CELLS IN THE BASE CASE.

By using the bidding actions computed using procedures 1 and 2, we have determined the evolution of the cumulative revenues as a function of time. The results are plotted on Fig. 5. As we can see, the cumulative revenues are not always growing. Indeed, they are decreasing during periods of time when the tank is filled with hydrogen. We note that at the end of the period 2007-2013, a cumulative revenue of 233,000€ is obtained.


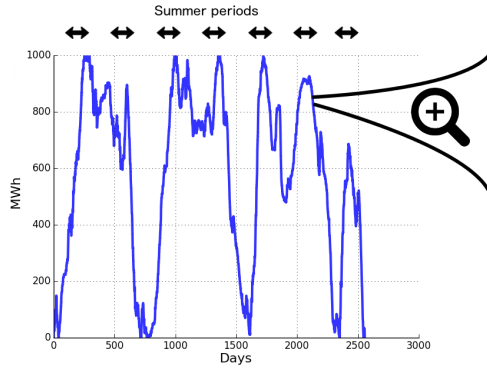



Fig. 5. cumulative revenues $\sum_t r(\bar{a}_t, p_t)$ as a function of time. Day 0 refers to the first of January 2007.

The evolution of the level of energy stored inside the storage tank s_t is shown on Fig. 6(a). It can be seen that hydrogen tends to be stored during summer and transformed back into electricity during winter. This is explained by the fact that most years, prices are higher in winter and lower in summer (see Fig. 7). Besides, daily fluctuations can also be



(a) Evolution throughout the years 2007 to 2013. Day 0 refers to the first January 2007.



(b) Zoom over two days

Fig. 6. Evolution of the energy reservoir level (s_t) as a function of time for the base case.


seen on Fig 6(b). Energy is accumulated during the night and transferred back to the grid during the day.

On Fig. 8, we have plotted the evolution of the price as a function of the period of the day. We can observe that with the years, the difference between on-peak and off-peak prices tends to decrease. More specifically, the peak prices occurring traditionally during the day tend to get much closer to the average price value. This can be explained by the significant investments that have been made after 2008 in photovoltaic panels. Let us now go back to Fig. 5 where we have plotted the evolution of the cumulative revenues over time. As one can observe, the rate of growth in cumulative revenues is higher for the first two years than for the rest of the period. This observation is a direct consequence from this flattening of the price evolution over the day.


b): Finally, we end this subsection by Fig. 9, which nicely illustrates on a single graphic the relation that exists behind the evolution of the prices and the sequence of actions taken.

B. Influence of the capacity of the storage tank on the maximum revenue

In this section, we study the revenues obtained as a function of the size of the reservoir. We have modeled the storage reservoir as varying between a few MWh up to a reservoir



(a) Average over all years



(b) Average over individual years

Fig. 7. Evolution of the average prices for the years 2007 to 2013 as a function of the period of the year.


which is large enough for never being fully filled by the agent. The results are plotted on Fig. 10. We remind the reader that in the previous subsection, a maximum capacity of 1000 MWh was used for the storage device. As we can see, the revenues are a growing function of the storage capacity. However, the incremental revenue obtained from the exploitation is lowered as the storage capacity increases. Whatever the size of the reservoir, it is not possible to generate a revenue which is larger than 272 000 €.

C. Influence of the discretization on the maximum revenue


In this section, we study the influence of the discretization steps δ_s and δ_u on the results obtained. To do so, we have run Procedure 1, followed by Procedure 2, for several values of δ_s and δ_u . Figure 11 plots the results obtained. Several interesting observations can be made. First, for a given value of δ_s (δ_u), the return of the bidding policy does not vary anymore when δ_u (δ_s) becomes lower than δ_s (δ_u). Second, if $\delta_u > \delta_s$ ($\delta_s > \delta_u$), better results can be obtained by moving δ_u closer to δ_s (δ_s closer to δ_u). Finally, in the case where the discretization steps are equal, the smaller they are, the better the quality of the policy. Note however, that below a certain value of the discretization steps, the quality of the policy remains roughly the same.

VII. REVENUES ESTIMATION UNDER MULTIPLE PRICE EVOLUTIONS

The experimental design exposed in the previous sections assumes that the future price evolution is known in advance. Additionally, it also relies on a discretization of the state-action space. A more realistic assumption would be to assume a set



(a) Average over all years



(b) Average over individual years

Fig. 8. Evolution of the price as a function of the hour of the day for the years 2007 to 2013.




Fig. 9. Illustration of the sequence of actions and of electricity prices with time. A dot is associated to every market period and its color refers to the action taken. Its position gives the market price at this time.

of possible price evolutions $\{(p_{k,t})_t\}_{k=1}^K$ where $K \in \mathbb{N}_0$ and consider the empirical return of a given sequence of actions over such a set of realizations as an approximation of the expected value of this sequence of actions.

If we make the hypothesis that the efficiencies of the fuel cell and the electrolyzer are constant, the reward function and the system dynamics are both linear mappings of states, actions and prices. The empirical return over K price outcomes is then also a linear mapping of the control variables. Finding a sequence of actions $(A_0^*, \dots, A_{n_D-1}^*)$ leading to the maximization of the empirical return can be achieved by solving




Fig. 10. Evolution of the expected revenues as a function of storage capacity for the years 2007 to 2013.




Fig. 11. Revenues generated by the bidding policies as a function of the action space discretization step (δ_u) and the state space discretization step (δ_s)

the following linear program:

$$\begin{aligned} & (A_0^*, \dots, A_{n_D-1}^*) \in \\ & \arg \max_{A_i \in \mathcal{A}_i, i=0 \dots n_D-1} \frac{1}{K} \sum_{k=1}^K \sum_{i=0}^{n_D-1} \rho(s_{i,0}, A_i, \vec{p}_{k,t}), \end{aligned}$$

where $\vec{p}_{k,t}$ denotes the vector of prices for the i -th day. Note that the previous linear program allows to solve the problem exposed in the previous section in the case $K = 1$ for continuous state-action spaces. The dynamic programming approach has the advantage (i) to ensure a linear complexity with the time horizon and (ii) to remain applicable in the case of non-linear and even non-convex system dynamics and reward function.

VIII. CONCLUSION

In this paper, a methodology has been proposed for estimating the revenues that can be generated by a high-capacity hydrogen-based storage device on the energy markets. It was

then used to estimate the revenues that could be generated on Belpex - the Belgian power exchange market.

The results show that for fixed size electrolyzers and fuel cells, significantly higher revenues can be achieved by having large storage capacities, such as for example hydrogen tanks that would take tens of days to fill or to empty. This is explained by the fact that with huge tanks, the storage device can be operated so as to exploit the inter-seasonal price fluctuations. The results also show that over the last years, the revenues that could have been generated by storage devices have decreased.

The research reported in this paper could be extended along several directions. First, our algorithm for estimating the future revenues assumes that the market price is not influenced by the storage device itself and, more importantly, that the future price evolution is known (or, at least, an ensemble of possible price evolutions are known). It would be worth extending the methodology proposed in this paper to a more general case. Note that this would imply working in a probabilistic setting where we would compute an expected future revenue or a distribution over future revenues.

Second, the only mechanism considered here for valorizing storage has been to buy or sell energy on the electricity market. But other mechanisms also exist, such as for example selling services to the balancing/reserves markets [18] or those that would relate to absorbing the excess of energy produced locally by renewable sources of energy so as to relieve congestions [19]. In this respect, it would be worth computing the revenues that can be generated by storage devices when all these mechanisms are taken into account.

Finally, it would be interesting to study how prediction models of the future revenues could be utilized to give clear indications about the storage technology in which to invest and about where to install storage devices.

Acknowledgments

The authors thank the "Consortium des Equipements de Calcul Intensif (CECI)", funded by the Fonds de la Recherche Scientifique de Belgique (F.R.S.-FNRS) under Grant No. 2.5020.11 for providing the computational resources needed for carrying out this research. Raphael Fonteneau is a post-doctoral fellow of the FRS-FNRS.

REFERENCES

- [1] CREG. Study (f)121011-cdc-1182 on "capacity remuneration mechanisms". Technical report, CREG, 2012.
- [2] Cristina Archer and Mark Jacobson. Supplying baseload power and reducing transmission requirements by interconnecting wind farms. *Journal of Applied Meteorology & Climatology*, 46(11), 2007.
- [3] Spyros Chatzivasileiadis, Damien Ernst, and Göran Andersson. The global grid. *Renewable Energy*, 57:372–383, 2013.
- [4] Spyros Chatzivasileiadis, Damien Ernst, and Göran Andersson. Global power grids for harnessing world renewable energy. *To be published by Elsevier*, 2014.
- [5] Nicola Armaroli and Vincenzo Balzani. The hydrogen issue. *ChemSusChem*, 4(1):21–36, 2011.
- [6] Jim Eyer and Garth Corey. Energy storage for the electricity grid: Benefits and market potential assessment guide. *Sandia National Laboratories*, 2010.

- [7] René Carmona and Michael Ludkovski. Valuation of energy storage: An optimal switching approach. *Quantitative Finance*, 10(4):359–374, 2010.
- [8] Nils Löhndorf and Stefan Minner. Optimal day-ahead trading and storage of renewable energies - an approximate dynamic programming approach. *Energy Systems*, 1(1):61–77, 2010.
- [9] Pedram Mokrian and Moff Stephen. A stochastic programming framework for the valuation of electricity storage. In *26th USAEE/IAEE North American Conference*, pages 24–27, 2006.
- [10] Stein-Erik Fleten and Trine Krogh Kristoffersen. Stochastic programming for optimizing bidding strategies of a nordic hydropower producer. *European Journal of Operational Research*, 181(2):916–928, 2007.
- [11] Daniel F Salas and Warren B Powell. Benchmarking a scalable approximate dynamic programming algorithm for stochastic control of multidimensional energy storage problems. Technical report, Working Paper, Department of Operations Research and Financial Engineering, Princeton, NJ, 2013.
- [12] Belpex, feb 2014.
- [13] Lucian Busoniu, Robert Babuska, Bart De Schutter, and Damien Ernst. *Reinforcement learning and dynamic programming using function approximators*. CRC Press, 2010.
- [14] Damien Ernst. *Near optimal closed-loop control Application to electric power systems*. PhD thesis, University of Liege, 2003.
- [15] Warren B Powell. What you should know about approximate dynamic programming. *Naval Research Logistics (NRL)*, 56(3):239–249, 2009.
- [16] Robert Dopp. Hydrogen generation via water electrolysis using highly efficient nanometal electrodes. *Quantum Sphere, Inc*, 2007.
- [17] Sivan Kartha and Patrick Grimes. Fuel cells: Energy conversion for the next century. *Physics Today*, 47(11):54–61, 2008.
- [18] Eurelectric. Decentralized storage: Impact on future distribution grids. Technical report, Eurelectric, 2012.
- [19] Quentin Gemine, Efthymios Karangelos, Damien Ernst, and Bertrand Cornélusse. Active network management: planning under uncertainty for exploiting load modulation. In *Bulk Power System Dynamics and Control-IX Optimization, Security and Control of the Emerging Power Grid (IREP), 2013 IREP Symposium*, pages 1–9. IEEE, 2013.