Delft University of Technology

Delft Center for Systems and Control

Technical report 20-003

Optimal control strategies for seasonal thermal energy storage systems with market interaction*

J. Lago, G. Suryanarayana, E. Sogancioglu, and B. De Schutter

If you want to cite this report, please use the following reference instead:
J. Lago, G. Suryanarayana, E. Sogancioglu, and B. De Schutter, "Optimal control strategies for seasonal thermal energy storage systems with market interaction," *IEEE Transactions on Control Systems Technology*, vol. 29, no. 5, pp. 1891–1906, Sept. 2021. doi:10.1109/TCST.2020.3016077

Delft Center for Systems and Control
Delft University of Technology
Mekelweg 2, 2628 CD Delft
The Netherlands
phone: +31-15-278.24.73 (secretary)
URL: https://www.dcsc.tudelft.nl

Optimal Control Strategies for Seasonal Thermal Energy Storage Systems with Market Interaction

Jesus Lago, Gowri Suryanarayana, Ecem Sogancioglu, and Bart De Schutter, *Fellow, IEEE*

***Abstract*—Seasonal thermal energy storage systems (STESSs) can shift the delivery of renewable energy sources and mitigate their uncertainty problems. However, to maximize the operational profit of STESSs and to ensure their long-term profitability, control strategies that allow them to trade on wholesale electricity markets are required. While control strategies for STESSs have been proposed before, none of them addressed electricity market interaction and trading. In particular, due to the seasonal nature of STESSs, accounting for the long-term uncertainty in electricity prices has been very challenging. In this paper, we develop the first control algorithms to control STESSs when interacting with different wholesale electricity markets. As different control solutions have different merits, we propose solutions based on model predictive control and solutions based on reinforcement learning. We show that this is critical since different markets require different control strategies: MPC strategies are better for day-ahead markets due to the flexibility of MPC whereas RL strategies are better for real-time markets because of fast computation times and better risk modeling. To study the proposed algorithms in a real-life setup, we consider a real STESS interacting with the day-ahead and imbalance markets in The Netherlands and Belgium. Based on the obtained results we show that: (1) the developed controllers successfully maximize the profits of STESSs due to market trading; (2) the developed control strategies make STESSs important players in the energy transition: by optimally controlling STESSs and reacting to imbalances, STESSs help to reduce grid imbalances.**

***Index Terms*—Optimal Control, Seasonal Storage Systems, Electricity Markets, Demand Response, Model Predictive Control, Reinforcement Learning.**

I. Introduction

While the energy transition [1] has the potential to highly improve our society, e.g. by mitigating climate change, it also poses some potential problems that need to be tackled [2]. Specifically, due to the weather dependence of renewable sources, a large integration of renewables implies more uncertain energy generation. In the case of electricity, as generation and consumption have to be balanced at all times, the more renewable sources are integrated, the more imbalances between generation and consumption occur, and the more complex the control and balance of the electrical grid becomes [3]. In this context, energy storage systems offer a promising solution for uncertain generation by providing flexibility and ancillary services, leading to smooth and reliable grid operation. [4].

A. Energy storage systems

Depending on the type of technology, there are different energy storage solutions [4], [5], e.g. lithium-ion batteries, pumped hydro storage, ultracapacitors, flywheels, molten-salt batteries, thermal storage systems, compressed air storage, or hydrogen storage. While most of these technologies can ensure efficient short-term and medium-term energy storage, efficient long-term energy storage has traditionally been more difficult to achieve: although some of these technologies can store energy for long periods, they are not economically very efficient [4]. However, long-term energy storage is arguably one of the most important elements to ensure the success of the energy transition. Particularly, as the share of wind and solar energy by 2030 is expected to reach very high levels (70-80% in some countries), and as the generation of renewables is seasonal dependent [5], seasonal energy storage solutions [5] that can store energy across several weeks or months are crucial in order to reduce seasonal fluctuations [4].

With regard to seasonal storage, there are primarily three solutions available that can provide electricity back to the grid: hydrogen storage, synthetic natural gas storage, and vanadium redox flow batteries [5], [6]. The first two approaches are power-to-gas technologies that make use of renewable sources to generate synthetic fuels, i.e. primarily hydrogen and methane [7]. The third belongs to the next-generation of batteries that can potentially store electricity for long horizon [6], [8]. In this context, besides vanadium redox flow batteries, there is also undergoing research into the next generation of post-lithium-ion technologies with capabilities of long-term storage [9], [10]. Despite their potential, these technologies still have several problems that make them economically non-viable: first, they are expensive technologies and in an early stage of development [6]–[12]. Second, synthetic fuels have a very low energy efficiency due to conversion losses [7]. Third, vanadium redox flow batteries and other post-lithium-ion batteries are yet not profitable and face multiple challenges that difficult their commercial deployment [6], [9], [10], [13].

Another option for storing energy over long horizons are *thermal energy storage (TES)* systems [14]. While in general

This research has received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 675318 (INCITE). The computational resources used in this work were provided by the VSC (Flemish Supercomputer Center), funded by the Research Foundation - Flanders (FWO) and the Flemish Government.

J. Lago is with the Delft Center for Systems and Control, Delft University of Technology, Delft, The Netherlands and with the Department of Algorithms, Modeling and Optimization, VITO-Energyville, Genk, Belgium (e-mail: j.lagogarcia@tudelft.nl)

G. Suryanarayana is with the Department of Algorithms, Modeling and Optimization, VITO-Energyville, Genk, Belgium (e-mail: gowri.suryanarayana@vito.be)

E. Sogancioglu is with the Department of Radiology and Nuclear Medicine, Radboud University, Nijmegen, The Netherlands (email: ecem.lago@radboudumc.nl)

B. De Schutter is with the Delft Center for Systems and Control, Delft University of Technology, Delft, The Netherlands (e-mail: b.deschutter@tudelft.nl)

these systems cannot provide electricity back to the grid, they are a more mature technology, have the advantage of being significantly less expensive than electrical energy storage [4], and can be used to satisfy heating and cooling demands.

In the context of TES technologies, there are three main categories: sensible heat storage, latent heat storage, and chemical energy storage [14], [15]. While the last two have higher energy densities, they are both more expensive and less mature, i.e. sometimes at the laboratory testing stage and with no large-scale seasonal project completed [14]. By contrast, sensible energy storage is the simplest, cheapest, most widespread, and most mature technology [15]. As a result, sensible heat storage systems are the focus of this paper. Note that, aligned with the literature [16]–[18], we use the name of *seasonal thermal energy storage systems (STESSs)* to refer to TES systems based on sensible heat storage.

B. Control of non-seasonal storage systems

The problem of controlling storage systems is a developed area of research that contains many approaches that consider market interaction. However, within this context, all the research has usually focused on short-term storage systems, i.e. non-seasonal storage. The aim of this section is to provide a brief overview of the different families of approaches within the field, describe which markets the control algorithms are designed for, and which control horizons are usually considered. It is important to note that, since the number of contributions to this field are numerous, this will not be a thorough literature review but a brief summary of the research field.

Optimization-based approaches have been employed in numerous applications [19]–[27] and are arguably the most widely used family. In order to interact with different markets, these approaches are formulated as sequential multi-stage optimization problems. Another family of approaches are based on dynamic programming and Markov processes [28]–[30]. While these approaches often provide global optimal solutions, do not scale for large systems [31]. A third family are rule-based approaches [23], [32], which derive a set of logical rules to control the storage systems. Finally, there are game-theoretical models [33], which are based on competition economic models.

In terms of markets, control approaches have been proposed for many different cases. The most common of them is trading in the day-ahead market together with the balancing market [19]–[21], [27], [33] or with the real-time market [24]–[26]. Other proposed strategies include: frequency regulation coupled with energy arbitrage markets [29]; day-ahead market [30]; primary frequency response market [32]; real-time markets [22]; or day-ahead, intraday, and balancing markets [23], [28]. To the best of our knowledge, approaches that exploit the imbalance markets have not been proposed.

In terms of the horizon, the majority of the approaches perform price arbitrage between day-ahead and markets closer to real-time considering optimization horizons of one day [19]–[29]. In this context, no approaches provide solutions for trading energy over long horizons, e.g. months.

C. Control of seasonal storage systems

In the context of seasonal storage systems, several optimal control strategies have been also proposed. However, none of the proposed methods are designed for market interaction. In [17] and [34], *model predictive control (MPC)* based strategies are proposed to control aquifer thermal energy storage systems; however, while the controller is designed to satisfy physical constraints and a stochastic heat demand, the STESS does not interact with electricity markets. Similarly, in [35], a dynamic programming approach is proposed to control borehole thermal storage systems; however, the controller assumes a constant market price and does not distinguish between different markets. In [18] and [36], two control algorithms are proposed to control solar communities with a borehole thermal storage system; however, similar to other studies, price and markets are not considered and the controller is limited to satisfy the system constraints and the heat demand. In [16], a data-driven stochastic predictive control scheme to operate an energy hub with seasonal storage capabilities is proposed; the goal of the approach is to minimize the total energy consumption and be cost efficient; however, here also, the algorithm does not consider real market prices nor market trading. Similarly, [37] proposes an optimal charging strategy for borehole thermal storage systems; however, the focus of the controller is to maximize the renewable energy use and to reduce CO_2 emissions, and also here, no prices nor market interaction are considered.

D. Motivation of the research

While the field of control for storage systems features several approaches, they are either limited to approaches for short-term storage with market interaction or seasonal storage without market interaction.

Generic methods for storage systems, while they model market interaction, cannot cope with long optimization horizons. Particularly, all the existing methods [19]–[21], [23], [27]–[29], [33] provide trading approaches where storage systems trade energy with daily/weekly horizons and use price differences to perform price arbitrage. This poses a challenge for seasonal storage systems like STESSs, where the optimization has to be perform over yearly horizons. The reason why the existing methods cannot be applied to STESSs is twofold:

1) STESSs require forecasts of electricity prices over yearly horizons. While there are several forecast methods [3], [38], [39] for short-term horizons, i.e. days, there are no reliable methods to forecasts for long-term horizons.
2) Because of the long optimization horizons, the number of variables in the optimization problems grows very large. In this context, the existing methods become computationally intractable, e.g. many of them are based on mixed-integer optimization.

In the context of control algorithms for seasonal storage, while long horizons are sometimes considered, none of the existing methods are able to model electricity market interaction. This interaction is of primary importance for several reasons:

1) To maximize the profit of STESSs, they should be allowed to interact with markets. In particular, while controlling STESSs to satisfy heat demand and/or to maximize renewable energy usage are important goals, they do not necessarily optimize the economic cost of STESSs. This is specially important to increase the number of storage systems in the electrical grid: if the time for return on investment of STESSs is too long, STESSs might become unattractive investments.
2) As we will show in this paper, the profits of the STESSs are maximized when interacting with multiple markets. Therefore, controlling STESSs based on a single price or a single market is economically suboptimal.
3) To help reduce grid imbalances, STESSs need to be able to arbitrage in more than one market. In particular, to provide up-regulation in the imbalance markets, i.e. a real-time market, STESSs need to first buy that electricity in a market with an earlier gate closure time.

E. Contributions

To fill the scientific gap described earlier, we present four contributions in this paper:

- We propose and develop different control strategies for STESSs that can interact with multiple wholesale electricity markets. In particular, considering that there are several trading markets for STESSs, we propose control approaches for two cases: interaction with the day-ahead market alone, and simultaneous interaction with the day-ahead and imbalance markets. In addition, as different control approaches have different merits, for each market interaction we propose an MPC-based controller and an RL-based controller.
- We propose the first control algorithms for storage systems that consider long optimization horizons. Particularly, unlike the existing literature on seasonal storage systems, the proposed methods quantify the price variations and uncertainty over a horizon of a year, and exploit these variations to maximize the profits of the storage system. In the case of the MPC approaches, this is obtained using a novel two-stage optimization problem, a forecasting method for long horizons, and a variable time grid formulation. In the case of the RL approaches, the solution involves a new simulation framework for long horizons and a collaborative RL strategy.
- We assess the merits of each control solution for the different markets and show that, while MPC-based methods are most suitable for day-ahead markets, RL-based methods perform better when trading in the imbalance market.
- Finally, we empirically demonstrate that STESSs can play an important role in the energy transition by helping grid operators to reduce grid imbalances. We show that the economic incentives of STESSs are aligned with the regulatory duties of the grid operators and that STESSs can help balancing the grid to allow further integration of renewable sources. To the best of our knowledge, this is the first time that trading on the imbalance market is explicitly evaluated from the perspective of balancing the grid and the regulatory duties of the TSO.

We also have two additional contributions: we propose a simple scenario generation method for generating long-term price scenarios and a novel method for imbalance price forecasting. This contribution refers specifically to forecasting imbalance prices and not real-time *local marginal prices (LMPs)*. Although for the latter there are already forecasting methods [40], [41], imbalance prices have different properties than real-time LMPs and are much harder to predict.

F. Organization of the paper

The paper is organized as follows: Section III introduces and defines the framework of a general STESS interacting with electricity markets. Sections IV and V present respectively the proposed MPC and RL approaches. Finally, Section VI studies the performance of the proposed control approaches under several case studies and considering a real STESS. Appendix A describes the proposed scenario generation method, Appendix B explains the imbalance price forecasting method, Appendix C introduces and defines wholesale electricity markets, and Appendix D presents the theoretical basis of MPC and RL.

II. Motivation for the Selected Methodology

Designing controllers for STESSs that trade in multiple electricity markets is a very challenging task as selecting the right control algorithms or right markets is not straightforward.

A. Control algorithms for STESSs with market trading

Considering the difficulty of market trading, state-of-the-art control approaches, e.g. MPC [42] or *reinforcement learning (RL)* [43], are highly desirable. However, in the case of MPC [42], several problems appear:

- MPC requires realistic forecasts and/or scenarios of electricity prices over yearly horizons. While there are several forecast methods [3], [38], [39] and scenario generation methods [44]–[46] for short-term horizons, i.e. days, there are no reliable methods, to the best of our knowledge, to forecasts or generate scenarios for long-term horizons.
- In real-time electricity markets, e.g. imbalance markets, an action has to be taken within seconds. As the MPC works with a year horizon and the price resolution is typically 15 min, the number of variables in the optimization problem grows very large. As a result, MPC suffers from computational tractability problems to provide the optimal action within the available time frame.

While data-driven and RL techniques can mitigate or solve these two issues, they also have problems of their own:

- While they do not require forecasts or scenarios of electricity prices, they need to generate artificial time series of electricity prices to simulate the market conditions. Thus, a method to generate realistic prices is still needed.
- As they are trained offline, they do not have the real-time computation issues of MPC. However, that comes at the cost of adaptability: if market conditions change or if the STESS suffers from a problem, e.g. a heat

exchanger breaks, the controller has to be re-trained again. As the training can take several days, this limits the adaptability of RL to changes in the environmental conditions. In contrast, as MPC computes the solution online, any change in the environment can be directly included as a change in the optimization problem or by re-estimating the dynamical model with little impact on computation cost.
- The solutions of RL are at best a good approximation of the optimal solution while MPC obtains an optimal solution by explicitly solving the given control problem.
- Unlike MPC, RL cannot explicitly model hard constraints (they can only be modeled as part of the reward). As such, RL cannot guarantee that the provided solutions do not violate constraints.

Based on these arguments, it becomes clear that the perfect method does not exist and considering RL or MPC involves several trade-offs. As a result, for this research, we will propose different methods based on the two families and analyze the performance of each.

B. Electricity markets for trading with STESSs

Another important point to consider is that not all electricity markets are the same. While in theory STESSs could trade in any electricity market, there are two trading strategies that are especially relevant: trading only in the day-ahead market and trading in both the day-ahead and the imbalance market. Trading only in the day-ahead market is arguably the safest trading strategy for STESSs as the day-ahead market is the electricity market with the largest volume of renewable energy trading, i.e. with low but volatile prices, and players incur no risks as they submit bidding curves.

While trading only in the day-ahead market is a low-risk and cost-effective trading strategy, it might still not be the most optimal economic strategy for STESSs. In particular, while on average, prices in the imbalance market are larger than in the day-ahead market, since the imbalance prices are much more volatile, there are periods of time where imbalance prices are much lower (sometimes becoming even negative). In addition, by participating in the imbalance market, STESSs might be able to help reduce grid imbalances: as during periods of positive imbalances, i.e. generation larger than consumption, prices are low, STESSs could wait for these periods to buy their energy; by doing so, they would not only reduce grid imbalances but also increase their own profits. Similarly, as prices are high during periods of negative imbalances, STESSs can make use of their charging flexibility to first buy energy in the day-ahead market, and then sell it in the imbalance market if imbalances are negative or use it if they are positive. By doing so, STESSs could potentially increase their profits while helping to reduce negative imbalances.

It is important to note that, despite all these potential benefits, trading strategies for the imbalance market have much higher risks: in the imbalance market, agents take an action for the next time interval without knowing the imbalance price. Particularly, as imbalance prices are based on the grid imbalances during a period of time, the price is only known after the period is over. Thus, trading strategies for the imbalance market heavily rely on price forecasters and have an associated risk.

In this paper, we will explore both trading strategies, i.e. trading in just the day-ahead market and trading in both the day-ahead and imbalance markets, and study the benefits of each.

III. Seasonal Storage System Framework

In order to introduce the control algorithms, we need to define the framework of a general STESS interacting with the electricity markets. For notational simplicity, concatenations of several vectors, e.g. $[x^\top, y^\top]^\top$, will be shortened as (x, y).

A. Dynamical model

An STESS can be defined as a general dynamical system with an internal state $\mathbf{x}(t)$, controls $\mathbf{u}(t) = (\dot{\mathbf{Q}}^{\mathrm{in}}(t), \dot{\mathbf{Q}}^{\mathrm{out}}(t))$, n_{units} storage units, and external disturbances $\mathbf{d}(t)$. The internal state $\mathbf{x}(t)$ represents the state of charge of the system. The controls $\dot{\mathbf{Q}}^{\mathrm{in}}(t) \in \mathbb{R}^{n_{\mathrm{in}}}$ and $\dot{\mathbf{Q}}^{\mathrm{out}}(t) \in \mathbb{R}^{n_{\mathrm{out}}}$ respectively represent the rate at which energy is injected and extracted into/from the system. The disturbance represents any uncontrollable input, e.g. the external temperature.

The dynamics of the system are defined by a *partial differential equation (PDE)*. For a sensible heat storage device with water stratification, the system can be divided into n_{units} layers acting as individual storage units, and the dynamics of a layer i represented by the following PDE [47]:

$$\frac{\partial x_i}{\partial t} = a_1 \frac{\partial^2 x_i}{\partial z^2} + a_2(d - x_i) + a_3(\dot{Q}_i^{\mathrm{in}} - \dot{Q}_i^{\mathrm{out}}), \qquad (1)$$

where z represents the direction of stratification.

B. Heat demand and purchased energy

In general, an STESS is required to supply an uncertain heat demand $\dot{Q}^{\mathrm{d}}(t)$. To do so, an STESS buys energy $\dot{Q}^{\mathrm{m}}(t)$ from some market, stores it, and then delivers it to follow $\dot{Q}^{\mathrm{d}}(t)$. To maximize the profits, it needs to consider the price of $\dot{Q}^{\mathrm{m}}(t)$, the storage efficiency, and an estimation of the future heat demand $\dot{Q}^{\mathrm{d}}(t)$. Therefore, the following holds:

$$\dot{Q}^{\mathrm{d}}(t) = \sum_{i=1}^{n_{\mathrm{out}}} \dot{Q}_i^{\mathrm{out}}(t), \qquad \dot{Q}^{\mathrm{m}}(t) = \sum_{i=1}^{n_{\mathrm{in}}} \dot{Q}_i^{\mathrm{in}}(t), \qquad (2)$$

i.e. the heat demand should equal the sum of the energy extracted from the STESS and the energy bought in the market should equal to sum of the energy introduced in the STESS.

C. Trading in the day-ahead market

Given a day-ahead market with unknown daily hourly prices $(\lambda_1^{\mathrm{dam}}, \ldots, \ldots \lambda_{24}^{\mathrm{dam}})$, the goal of any control algorithm for an STESS is to build optimal bidding curves to maximize the profit. In particular, the aim is to, one day in advance, build 24 optimal bidding curves $\dot{Q}_1^{\mathrm{b}}(\cdot), \ldots, \dot{Q}_{24}^{\mathrm{b}}(\cdot)$ such that, while the STESS always has enough energy to satisfy the demand $\dot{Q}^{\mathrm{d}}(t)$, the cost of the purchased power $\dot{Q}^{\mathrm{dam}}(t)$ is minimized. In this market structure, the purchased power $\dot{Q}^{\mathrm{dam}}(t)$ at every hour h is defined by:

$$\dot{Q}^{\mathrm{dam}}(t) = \dot{Q}_h^{\mathrm{b}}(\lambda_h^{\mathrm{dam}}), \quad \forall\, t \in [h, h+1]. \qquad (3)$$

D. Trading in the imbalance market

For the imbalance market, the imbalance price λ^{imb} is always unknown when purchasing/selling power as the price λ^{imb} is determined in real time by the reserves activated by the TSO. In particular, at time step k, a market agent has to decide whether to sell, buy, or not trade without knowing the imbalance price λ_k^{imb} for the interval. As λ_k^{imb} is usually known immediately at the next interval, the agent can take the decision based on past imbalance prices $\lambda_{k-1}^{\mathrm{imb}}, \lambda_{k-2}^{\mathrm{imb}}, \ldots$ or any other information available at time step $k-1$.

Defining as $\dot{Q}^{\mathrm{imb}}(t)$ the energy traded in the imbalance market, with positive and negative values respectively representing energy that is bought and sold, it holds that:

$$-\dot{Q}^{\mathrm{imb}}(t) \leq \dot{Q}^{\mathrm{dam}}(t), \tag{4}$$

i.e. the energy sold in the imbalance market by an STESS is limited by the energy purchased on any previous market (the day-ahead market in the case of the proposed control algorithms). Particularly, because the STESS cannot effectively convert heat back to electricity, any energy sold is limited by the energy bought within the same day in other markets, and the STESS cannot sell any energy that was previously stored. Similarly, it holds that:

$$\dot{Q}^{\mathrm{m}}(t) = \dot{Q}^{\mathrm{dam}}(t) + \dot{Q}^{\mathrm{imb}}(t), \tag{5}$$

i.e. the total energy purchased for the STESS is the sum of the energy purchased in the day-ahead and imbalance markets.

Considering these definitions, a control algorithm for the imbalance market has to select the value of $\dot{Q}^{\mathrm{imb}}(t)$ for each time step k so that, while the STESS has enough energy to to satisfy the demand $\dot{Q}^{\mathrm{d}}(t)$, the total cost of trading $\dot{Q}^{\mathrm{dam}}(t)$ and $\dot{Q}^{\mathrm{imb}}(t)$ is minimized. To do so, the control algorithm receives as an input the energy $\dot{Q}^{\mathrm{dam}}(t)$ purchased in the day-ahead, and selects the value of $\dot{Q}^{\mathrm{imb}}(t)$.

IV. MPC APPROACHES

In this section, we derive and explain the two proposed MPC approaches: one to trade exclusively on the day-ahead market, and a second one to trade on both the day-ahead and the imbalance market.

A. Bidding functions

In the case of the day-ahead electricity market, the goal of the MPC is to provide the 24 optimal bidding functions $\dot{Q}_h^{\mathrm{b}}(\cdot)$, for $h = 1, 2, \ldots, 24$. Since standard MPC can only provide the optimal market power $\dot{Q}_{\tilde{\lambda}}^{\mathrm{dam}}$ for a fixed price $\tilde{\lambda}$, an additional step is needed. For each hour h, we propose the following approach:

1) Predefine n_{p} discrete prices $\{\lambda^1, \lambda^2, \ldots, \lambda^{n_{\mathrm{p}}}\}$ for the price λ^{dam} at hour h.
2) Fix the remaining 23 day-ahead prices using their expected value, e.g. a forecast.
3) Solve the MPC for each of these n_{p} prices and obtain the associated optimal market powers $\{\dot{Q}_{\lambda^1}^{\mathrm{dam}}, \dot{Q}_{\lambda^2}^{\mathrm{dam}}, \ldots, \dot{Q}_{\lambda^{n_{\mathrm{p}}}}^{\mathrm{dam}}\}$ at hour h.
4) Build the bidding function as a piecewise constant function based on the obtained solutions:

$$\dot{Q}_h^{\mathrm{b}}(\lambda^{\mathrm{dam}}) = \begin{cases} \dot{Q}_{\lambda^1}^{\mathrm{dam}}, & \lambda^{\mathrm{dam}} \leq \lambda^1 \\ \dot{Q}_{\lambda^2}^{\mathrm{dam}}, & \lambda^1 < \lambda^{\mathrm{dam}} \leq \lambda^2 \\ \vdots & \\ \dot{Q}_{\lambda^{n_{\mathrm{p}}}}^{\mathrm{dam}}, & \lambda^{n_{\mathrm{p}}-1} < \lambda^{\mathrm{dam}} \leq \lambda^{n_{\mathrm{p}}} \\ 0, & \lambda^{n_{\mathrm{p}}} < \lambda^{\mathrm{dam}} \end{cases} \tag{6}$$

This approach for building bidding functions is obviously only possible as long as the bidding functions within one day are independent of each other. However, since STESSs are very large storage devices, their internal state does not vary much within one day. As a result, the choice of one bidding function does not affect much the others and the assumption of independent bidding functions holds in practice.

Moreover, due to the market structure and the long optimization horizons of STESS, the 24 bidding functions are very similar. In detail, as the 24 daily bids are submitted at the same time, all the bids are built based on the same information, e.g. the STESS state. Moreover, as the STESS is flexible, it does not matter at which hour of the day it buys energy: because of the large storage size of the STESS, the state of the STESS barely changes with the action taken in a given hour. As such, the STESS states between consecutive days never differ too much, and, as the optimal bidding functions only depend on the STESS state, it follows that the optimal bidding function for every hour of a given day are similar. As a result, in a given day, the difference in price distribution between hours is not important, and the STESS reacts almost equally to a market price independently of the hour, i.e.:

$$\dot{Q}_1^{\mathrm{b}}(\lambda^{\mathrm{dam}}) \approx \dot{Q}_2^{\mathrm{b}}(\lambda^{\mathrm{dam}}) \approx \ldots \approx \dot{Q}_{24}^{\mathrm{b}}(\lambda^{\mathrm{dam}}), \quad \forall\ \lambda^{\mathrm{dam}}. \tag{7}$$

Thus, to build the 24 bidding functions, it is only needed to obtain the bidding function $\dot{Q}_1^{\mathrm{b}}(\cdot)$ for the first hour.

B. MPC for day-ahead trading

As motivated in the previous section, we only need to estimate the bidding function $\dot{Q}_1^{\mathrm{b}}(\cdot)$ for the first hour of the day. However, instead of solving a single OCP like in standard MPC, we need to discretize the first price λ_1^{dam} into a discrete set of prices $\{\lambda^1, \lambda^2, \ldots, \lambda^{n_{\mathrm{p}}}\}$, and for each of these prices solve the relevant OCP.

For the sake of simplicity, in this section we will assume that each OCP is optimized using a discrete time grid $t_1, t_2, \ldots, t_{N+1}$, i.e. using an optimization horizon equal to $t_{N+1} - t_1$; the details of how the time grid is defined will be covered in Section IV-D. Similarly, we will assume that the expected day-ahead prices $\{\bar{\lambda}_k^{\mathrm{dam}}\}_{k=1}^N$, the expected heat demand values $\{\bar{\dot{Q}}_k^{\mathrm{d}}\}_{k=1}^N$, and the expected disturbances $\{\bar{\mathbf{d}}_k\}_{k=1}^N$ are also provided; the method to obtain these values are explained in Appendix A.

Considering the previous definitions, at every day and for each price λ^j, the MPC approach solves the following OCP:

OCP(λ^j):

$$\underset{\substack{\mathbf{x}_1,\dot{\mathbf{Q}}_1^{\text{in}},\dot{\mathbf{Q}}_1^{\text{out}},\dot{Q}_1^{\text{dam}},\mathbf{x}_2,\ldots,\\ \dot{\mathbf{Q}}_N^{\text{in}},\dot{\mathbf{Q}}_N^{\text{out}},\dot{Q}_N^{\text{dam}},\mathbf{x}_{N+1}}}{\text{minimize}} \quad \lambda^j\,\dot{Q}_1^{\text{dam}} + \sum_{k=2}^{N} \bar{\lambda}_k^{\text{dam}}\,\dot{Q}_k^{\text{dam}} \tag{8a}$$

subject to

$$\mathbf{x}_1 = \tilde{\mathbf{x}}_1, \tag{8b}$$

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k, \dot{\mathbf{Q}}_k^{\text{in}}, \dot{\mathbf{Q}}_k^{\text{out}}, \bar{\mathbf{d}}_k), \quad \text{for } k = 1,\ldots N, \tag{8c}$$

$$\dot{Q}_k^{\text{dam}} \leq \dot{Q}_{\max}^{\text{m}}, \quad \text{for } k = 1,\ldots N, \tag{8d}$$

$$\sum_{i=1}^{n_{\text{in}}} \dot{Q}_{k,i}^{\text{in}} = \dot{Q}_k^{\text{dam}}, \quad \text{for } k = 1,\ldots N, \tag{8e}$$

$$\sum_{i=1}^{n_{\text{out}}} \dot{Q}_{k,i}^{\text{out}} = \bar{\dot{Q}}_k^{\text{d}}, \quad \text{for } k = 1,\ldots N, \tag{8f}$$

$$0 \leq \dot{\mathbf{Q}}_k^{\text{in}} \leq g_{\text{in}}(\mathbf{x}_k), \quad \text{for } k = 1,\ldots N, \tag{8g}$$

$$0 \leq \dot{\mathbf{Q}}_k^{\text{out}} \leq g_{\text{out}}(\mathbf{x}_k), \quad \text{for } k = 1,\ldots N, \tag{8h}$$

$$\mathbf{x}_{\min} \leq \mathbf{x}_k \leq \mathbf{x}_{\max} \quad \text{for } k = 1,\ldots N, \tag{8i}$$

$$\mathbf{x}_{N+1} = \tilde{\mathbf{x}}_1, \tag{8j}$$

where:

- The objective function represents the cost of purchasing energy considering that the first price is fixed and given by λ^j and that the remaining prices in the horizon are the expected prices in the market.
- Equation (8b) fixes the initial state, which is assumed to be known and given by $\tilde{\mathbf{x}}_1$.
- Equation (8c) ensures that the dynamics of the system are ensured at every time step. To discretize the continuous PDE, i.e. (1), we consider an explicit Euler integration scheme [47] as it provides a good trade-off between speed and accuracy for long optimization horizons
- To model the discrete dynamics, a multiple shooting [48] scheme is used. Unlike single shooting, multiple shooting explicitly includes the state x in the optimization problem. This is done to obtain a sparse Hessian and an easier to optimize problem.
- The maximum power purchased from the market is limited by (8d).
- Equation (8e) ensures that the input power equals the power purchased from the market.
- Through (8f) it is ensured that the heat demand is met.
- Equations (8g) and (8h) ensure the individual charging and discharging limits of each individual storage device. The upper limit is usually a function of the state as the maximum power that can be charged/discharged usually depends on the state of charge.
- The limits on the STESS state are defined by (8i).
- The OCP should avoid depleting the STESS at the end of the horizon. To do so, as the optimization horizon is usually a seasonal periodic cycle (see Section IV-D for details), (8j) constrains the STESS to have the same state of charge at the beginning and at the end.
- The objective function is simplified to leave out some costs, e.g. maintenance costs, start-up costs, or utility costs. Simplifying the objective function to only include the market cost is a design choice motivated by two reasons: first, some of these costs are orders of magnitude lower than the market cost[1]. Second, some costs simply offset the profitability by a constant or a scaling factor and are not relevant for the control algorithm.

After solving an OCP for each discrete price λ^{j}, the optimal bidding function $\dot{Q}_1^{\text{b}}(\cdot)$ can be estimated using (6), where the optimal market power $\dot{Q}_{\lambda^{\mathrm{j}}}^{\text{dam}}$ equals $\dot{Q}_1^{\text{dam}}$.

C. MPC for day-ahead and imbalance trading

The MPC-based approach to trade in both the day-ahead and the imbalance market consists of two separate MPC algorithms that run one after the other:

- A first MPC algorithm that trades in the day-ahead market but, unlike the MPC algorithm defined in the previous section, it considers that there is also a possible future interaction with the imbalance market.
- A second MPC algorithm that trades in the imbalance market and that considers that there is also possible future interactions with the day-ahead market. However, unlike the MPC algorithm for the day-ahead market, it runs on real time and it does not build bidding functions. Instead, at time step $k-1$, it considers a forecast $\hat{\lambda}_k^{\text{imb}}$ of the next imbalance price and then solves a single OCP to obtain the optimal power $\dot{Q}_k^{\text{imb}}$ to trade in the imbalance market.

It is important to note that, as with the day-ahead market, both algorithms are based on deterministic MPC. Given the uncertainty in electricity prices, one could argue that a more optimal approach would be to employ stochastic MPC. However, due to the long horizons involved, the computation time required for stochastic MPC makes the approach infeasible for real-time application (especially for the imbalance market). Particularly, for trading in the imbalance market, the MPC approach already requires (in the deterministic setting) to approximate the 1-year horizon to 1 month, i.e. the obtained optimal solution is approximated and no longer optimal w.r.t. the yearly seasonal period. A stochastic setting would only make this approximation worse. While larger computation capabilities could perhaps mitigate the issue, there are is another problem: as MPC solves a non-convex problem, there is no guarantee on the maximum computation time and more computational power might not help much.

C.1 MPC for the day-ahead market

To define the OCP of the first MPC algorithm, we will again consider that the discrete time grid $t_1, t_2, \ldots, t_{N+1}$, the expected day-ahead prices $\{\bar{\lambda}_k^{\text{dam}}\}_{k=1}^N$, imbalance prices $\{\bar{\lambda}_k^{\text{imb}}\}_{k=1}^N$, heat demand values $\{\bar{\dot{Q}}_k^{\text{d}}\}_{k=1}^N$, and disturbances $\{\bar{\mathbf{d}}_k\}_{k=1}^N$ are given. In addition, we will simplify the vector of input controls by $\mathbf{u}_k = (\dot{\mathbf{Q}}_k^{\text{in}}, \dot{\mathbf{Q}}_k^{\text{out}}, \dot{Q}_k^{\text{dam}}, \dot{Q}_k^{\text{imb}})$. Then, at every day and for each discrete price in $\{\lambda^1, \lambda^2, \ldots, \lambda^{n_{\mathrm{p}}}\}$, the MPC solves the following OCP:

[1] This information was obtained from the case study site company.

OCP(λ^j):

$$\underset{\substack{\mathbf{x}_1,\mathbf{u}_1,\mathbf{x}_2,\ldots,\\ \mathbf{u}_N,\mathbf{x}_{N+1}}}{\text{minimize}} \quad \lambda^j\, \dot{Q}_1^{\text{dam}} + \sum_{k=2}^{N} \bar{\lambda}_k^{\text{dam}}\, \dot{Q}_k^{\text{dam}} + \sum_{k=1}^{N} \bar{\lambda}_k^{\text{imb}}\, \dot{Q}_k^{\text{imb}} \tag{9a}$$

subject to

$$\mathbf{x}_1 = \tilde{\mathbf{x}}_1, \tag{9b}$$

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k, \dot{\mathbf{Q}}_k^{\text{in}}, \dot{\mathbf{Q}}_k^{\text{out}}, \bar{\mathbf{d}}_k), \quad \text{for } k = 1, \ldots N, \tag{9c}$$

$$\dot{Q}_k^{\text{dam}} + \dot{Q}_k^{\text{imb}} \leq \dot{Q}_{\max}^{\text{m}}, \quad \text{for } k = 1, \ldots N, \tag{9d}$$

$$\sum_{i=1}^{n_{\text{in}}} \dot{Q}_{k,i}^{\text{in}} = \dot{Q}_k^{\text{dam}} + \dot{Q}_k^{\text{imb}} \quad \text{for } k = 1, \ldots N, \tag{9e}$$

$$\sum_{i=1}^{n_{\text{out}}} \dot{Q}_{k,i}^{\text{out}} = \bar{\dot{Q}}_k^{\text{d}}, \quad \text{for } k = 1, \ldots N, \tag{9f}$$

$$0 \leq \dot{\mathbf{Q}}_k^{\text{in}} \leq g_{\text{in}}(\mathbf{x}_k), \quad \text{for } k = 1, \ldots N, \tag{9g}$$

$$0 \leq \dot{\mathbf{Q}}_k^{\text{out}} \leq g_{\text{out}}(\mathbf{x}_k), \quad \text{for } k = 1, \ldots N, \tag{9h}$$

$$\mathbf{x}_{\min} \leq \mathbf{x}_k \leq \mathbf{x}_{\max} \quad \text{for } k = 1, \ldots N, \tag{9i}$$

$$\dot{Q}_k^{\text{dam}} \geq 0, \quad \text{for } k = 1, \ldots N, \tag{9j}$$

$$-\dot{Q}_k^{\text{dam}} \leq \dot{Q}_k^{\text{imb}}, \quad \text{for } k = 1, \ldots N, \tag{9k}$$

$$\mathbf{x}_{N+1} = \tilde{\mathbf{x}}_1. \tag{9l}$$

While the main structure is very similar to (8), there are some important differences:

- The algorithm minimizes the cost of purchasing energy like in (8a) but includes the future transactions in the imbalance market.
- The constraints that contain the power purchased from the market, i.e. (9d) and (9e), consider now the sum of the power purchased in both markets.
- Unlike the case of trading only in the day-ahead market, the STESS can now sell energy on the imbalance market if it has previously bought it in the day-ahead market. This is modeled by (9j) and (9k), which respectively guarantee that in the day-ahead market energy can only be bought, and that the energy sold in the imbalance market is limited to the energy bought in the day-ahead market.
- The amount of energy traded is not limited by the system demand. In particular, the total energy traded is limited by $\dot{Q}_{\max}^{\text{m}}$, which represents a safety upper bound that can be much larger than the heat demand $\dot{Q}^{\text{d}}$ and that simply models how risk-averse the STESS is to price arbitration.

C.2 MPC for the imbalance market

To define the second MPC algorithm, let us first make the following assumptions and definitions:

- The MPC algorithm for the imbalance market considers a new time grid $t'_1, t'_2, \ldots, t'_{N_1+1}$ with $t'_{N_1+1} \leq t_{N+1}$, i.e. a shorter horizon and a different discretization. The details on this discretization are provided in Section IV-D.
- The expected day-ahead prices $\{\bar{\lambda}_k^{\text{dam}}\}_{k=1}^{N_1}$, imbalance prices $\{\bar{\lambda}_k^{\text{imb}}\}_{k=1}^{N_1}$, heat demand values $\{\bar{\dot{Q}}_k^{\text{d}}\}_{k=1}^{N_1}$, and disturbances $\{\bar{\mathbf{d}}_k\}_{k=1}^{N_1}$ are again provided. (See Appendix A for details).
- The optimal state at time t_{N_1+1} is defined by $\mathbf{x}^\star_{N_1+1}$ and obtained from the solution of the MPC for the day-ahead market. In particular, this value can be obtained from the optimal solution of any of the n_{p} OCPs solved in the latest day-ahead market.
- An accurate forecast $\hat{\lambda}_1^{\text{imb}}$ of the next price in the imbalance market is available. The details of this forecast are explained in Appendix B.

Based on these definitions, before each imbalance market clearance, MPC solves the following OCP and trades the optimal solution $\dot{Q}_1^{\text{imb}}$ in the imbalance market:

$$\underset{\substack{\mathbf{x}_1,\mathbf{u}_1,\mathbf{x}_2,\ldots,\\ \mathbf{u}_{N_1},\mathbf{x}_{N_1+1}}}{\text{minimize}} \quad \hat{\lambda}_1^{\text{imb}}\, \dot{Q}_1^{\text{imb}} + \sum_{k=1}^{N_1} \bar{\lambda}_k^{\text{dam}}\, \dot{Q}_k^{\text{dam}} + \sum_{k=2}^{N_1} \bar{\lambda}_k^{\text{imb}}\, \dot{Q}_k^{\text{imb}} \tag{10a}$$

subject to

$$(9\text{b}) - (9\text{k}), \tag{10b}$$

$$\mathbf{x}_{N_1+1} = \mathbf{x}^\star_{N_1+1}. \tag{10c}$$

The new MPC scheme is very similar to the previous MPC for the day-ahead market but with some differences:

- As a bidding function is not needed, instead of solving the OCP multiple times for different possible prices, this MPC algorithm solves a single OCP considering the most likely imbalance price $\hat{\lambda}_1^{\text{imb}}$ in the next market clearance. Then, it trades directly the optimal solution $\dot{Q}_1^{\text{imb}}$ in the imbalance market.
- A distinction is made between the future expected imbalance prices $\{\bar{\lambda}_k^{\text{imb}}\}_{k=2}^{N_1}$ and the forecast price $\hat{\lambda}_1^{\text{imb}}$ in the next time step. This distinction is made because the accuracy of the forecast is better than that of the method used to generate the expected future values.
- As this MPC algorithm runs in real time, the computation time should be as small as possible. To reduce the computation time, a smaller horizon $t'_{N_1+1} < t_{N+1}$ is considered.
- As the optimization horizon t'_{N_1+1} is now smaller than a periodical seasonal cycle, it is suboptimal to constrain the final state to be equal to the initial state. However, not constraining the final state leads to an OCP that does not account for what happens after t'_{N_1+1}. To solve this problem, (10c) constrains the final state to be equal to the optimal state $\mathbf{x}^\star_{N_1+1}$ at time t'_{N_1+1}, which is obtained from the solution of the latest day-ahead MPC algorithm.

D. Time grid and optimization horizon

In the previous sections, we assumed that the discrete time grids where the OCPs were defined were given. In this section, we explain the methodology to define these time grids.

In general, to define a discrete-time grid, we also need to define the optimization horizon T and the discrete time step Δt. For an STESS, T represents its seasonal horizon, which is typically a year. While most applications consider a constant Δt along the optimization horizon, we argue that for an STESS this not necessary and should in fact be avoided:

- As day-ahead markets have a different price every hour, the minimum time step at the beginning of the horizon

is $\Delta t = 1\,\mathrm{h}$. However, due to the long optimization horizons, it is not possible to accurately estimate with hourly resolutions the price and demand distributions at the end of the horizon. Instead, it is better to estimate the distributions over larger intervals, e.g. several hours, where due to noise averaging the uncertainty can be better quantified.
- Another reason to consider a variable Δt is the computational cost: by increasing Δt towards the end of the horizon, we reduce the number of optimization points N and the computational complexity of the OCP.
- As MPC only needs the optimal control at the first time point, it can be argued that lowering the time resolution at the end of the horizon has little impact on the first optimal control.

Finally, based on T and Δt, the number of time intervals N is defined.

D.1 Day-head market

Considering that the day-ahead electricity market is cleared every day, the hourly resolution should only be needed for the first day. Based on this and the arguments above, for the day-ahead market MPC we consider a time grid $t_1, t_2, \ldots, t_{N+1}$ with a year horizon, using four different Δt, and containing $N = 1233$ time intervals:

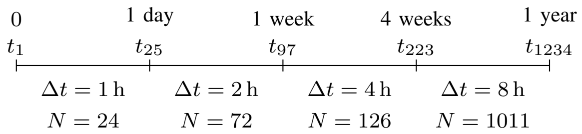

D.2 Imbalance market

For the case of the imbalance market, the minimum Δt is 15 minutes. Moreover, considering the large uncertainty in imbalances prices, we argue that the 15 min resolution is only needed for the first hour. Finally, as the MPC algorithm for the imbalance market runs in real time, the computation time should be as small as possible. Based on these arguments, we consider a time grid $t'_1, t'_2, \ldots, t'_{N_1+1}$ for the imbalance market with a horizon of four weeks, using four different Δt, and containing $N_1 = 225$ time intervals:

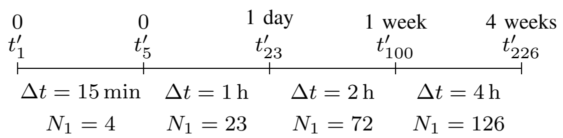

It could be argued that considering a horizon of four weeks instead of a year (the standard seasonal cycle) leads to suboptimal solutions, i.e. the MPC cannot account for what happens during a full seasonal cycle. However, as explained in Section IV-C2, MPC avoids this by constraining the state at the end of the four weeks to be equal to the optimal state $\mathbf{x}^\star_{226}$ at that time point.

V. RL APPROACHES

In this section, we present the two proposed RL approaches: one to trade in the day-ahead market, and a second one to trade in both the day-ahead and the imbalance markets.

A. RL for day-ahead trading

As with MPC, any RL control algorithm for the day-ahead market needs to estimate the bidding functions $\dot{Q}^{\mathrm{b}}_h(\cdot)$, for $h = 1, 2, \ldots, 24$. While in the case of MPC that required discretizing prices and solving multiple OCPs, for RL the bidding functions can be directly obtained from the optimal policy $\pi^\star(\mathbf{s}_k)$. In detail, if the RL agent is set up so that:

- The reward represents the cost of purchasing energy.
- The RL state $\mathbf{s}$ contains the day-ahead price λ^{dam}.
- The action $\mathbf{u}$ includes the power $\dot{Q}^{\mathrm{dam}}$ purchased from the market.

Then, by definition, the bidding function $\dot{Q}^{\mathrm{b}}(\lambda^{\mathrm{dam}})$ is implicitly defined by the optimal policy $\mathbf{u}^\star = \pi^\star(\mathbf{s}) = \pi^\star(\lambda^{\mathrm{dam}}, \ldots)$. Below we provide further details on the proposed RL algorithm.

A.1 State and control spaces

The first step to define the RL algorithm is to define its state and control spaces. For the proposed algorithm, the state $\mathbf{s} = (\mathbf{x}, \tau, \lambda^{\mathrm{dam}})$ is defined by three different features:

1) The state $\mathbf{x}$ of the STESS.
2) The time position τ within the periodic seasonal cycle, e.g. the day of the year.
3) The market price λ^{dam}.

The reason for selecting these three features is twofold:

- The optimal action $\mathbf{u}^\star = \pi^\star(\mathbf{s})$ can be selected based on both the state of the STESS and the environment.
- As we will show in Section V-A6, given a fixed time point $\tilde{\tau}$ and STESS state $\tilde{\mathbf{x}}$, the bidding function $\dot{Q}^{\mathrm{b}}(\lambda^{\mathrm{dam}})$ is by definition given by the optimal policy $\tilde{\pi}^\star(\tilde{\mathbf{x}}, \tilde{\tau}, \lambda^{\mathrm{dam}})$.

To define the action space $\mathbb{U}$, we consider that a single action $u \in \mathbb{R}^{n_{\mathrm{in}}+1}$ has the following format:

$$\mathbf{u} = (u_1, u_2, \ldots, u_{n_{\mathrm{in}}}, j). \tag{11}$$

In detail, we consider that each input control u_i can take $n_{\mathrm{dis}}+1$ discrete values uniformly separated between 0 and 1 and that the real power $\dot{Q}^{\mathrm{in}}_i$ into the storage i is obtained by multiplying u_i by the maximum power $\dot{Q}^{\max}_i$, i.e. $\dot{Q}^{\mathrm{in}}_i = u_i\,\dot{Q}^{\max}_i$. This scaling is done because $\dot{Q}^{\max}_i$ might depend on the system state and can change throughout time. For the output control, a single storage unit j is selected to provide the demand $\dot{Q}^{\mathrm{d}}$, i.e. $\dot{Q}^{\mathrm{d}} = \dot{Q}^{\mathrm{out}}_j$. The action space is then defined by the possible combinations of all these values.

A.2 Reward function

The reward r_k at time step k is defined as the negative of the cost of the energy purchased. Thus, assuming that the agent is at state $\mathbf{s}_k = (\mathbf{x}_k, \tau_k, \lambda^{\mathrm{dam}}_k)$ and takes an action $\mathbf{u}_k = (u_{1,k}, \ldots, u_{n_{\mathrm{in}},k}, j)$, r_k is defined as $-\lambda^{\mathrm{dam}}_k \sum_{i=1}^{n_{\mathrm{in}}} (u_{i,k} \cdot \dot{Q}^{\max}_{i,k})$. In addition, if the agent depletes the system and the demand $\dot{Q}^{\mathrm{d}}_k$ cannot be satisfied, the reward penalizes this situation with a cost 10 times larger than the cost of instantaneously buying $\dot{Q}^{\rm d}_k$ in the market[2]. Finally, as with standard RL algorithms, the reward at the last point in an episode is 0:

$$r_k = \begin{cases} 0, & \text{If } k = T_{\rm e} \\ -\lambda^{\rm dam}_k \left(\sum_{i=1}^{n_{\rm in}} (u_{i,k}\, \dot{Q}^{\rm max}_{i,k}) + 10\, \dot{Q}^{\rm d}_k \right), & \text{If the system is depleted} \\ -\lambda^{\rm dam}_k \sum_{i=1}^{n_{\rm in}} (u_{i,k} \cdot \dot{Q}^{\rm max}_{i,k}), & \text{Otherwise.} \end{cases} \tag{12}$$

A.3 Episode length and time grid

Another critical point when designing an RL algorithm is to select the episode length $T_{\rm e}$. For STESS, it can be argued that, to avoid optimal policies that deplete the STESS, the minimum $T_{\rm e}$ should be two seasonal periodic cycles. In particular, if the episode length equals the cycle length, the agent would know the time position within an episode as the agent knows the time position τ within a seasonal cycle. Using that information, the agent could potentially deplete the STESS at the end of the episode/cycle to reduce the cost. This behavior would be undesirable as the STESS needs to provide energy for more than a seasonal periodical cycle.

For the size of the discrete time grid, we consider that a time transition $k \rightarrow k+1$ spans a day. In particular, as with MPC, it is assumed that the state of charge does not change dramatically from one day to another and that the optimal bidding curves within a day are very similar. It is important to note that selecting this time step size is just a design choice and that it is equally possible to consider time steps of one hour at the expense of increasing the computation load.

A.4 Simulation environment

To train a RL agent to control STESSs, we use a simulation environment that recreates the world an STESS lives in. In detail, this environment consists of two modules:

- **STESS simulator**: a simulator of the dynamical model of the STESS: $\mathbf{x}_{k+1} = f(\mathbf{x}_k, \dot{\mathbf{Q}}^{\rm in}_k, \dot{\mathbf{Q}}^{\rm out}_k, \bar{\mathbf{d}}_k)$.
- **Environment simulator**: a simulator that produces realistic day-ahead market prices $\lambda^{\rm dam}$, heat demand $\dot{Q}^{\rm d}$, and disturbances $\mathbf{d}$. To obtain a simulator that generates realistic time series, the method for scenario generation explained in Appendix A is considered.

A.5 Training algorithm

The last step before training the agent is to select the specific RL algorithm to estimate the optimal policy $\pi^\star(\mathbf{s})$. For the case of STESSs, we propose using fitted Q-iteration [43], [49] with boosting trees [50]. The reason for selecting this algorithm is that we empirically observed (using the real system presented in Section VI) that this algorithm performed as good as more advanced RL algorithms but without the additional computational complexity. Unlike the deterministic MPC approach, price uncertainty is implicitly included in this approach as the RL agent is trained with a probabilistic reward. Therefore, the RL agent can learn some notion of risk that quantifies the distribution of a reward for a given state.

[2]Selecting a factor of 10 is a design choice. The agent just needs a large penalty cost whenever it depletes the STESS.

A.6 Building bidding functions

After the RL agent is trained, the optimal bidding functions $\dot{Q}^{\rm b}(\cdot)$ are directly obtained. In particular, given a fixed time point $\tilde{\tau}$ and STESS state $\tilde{\mathbf{x}}$, we have an optimal policy $\mathbf{u}^\star = \pi^\star(\tilde{\mathbf{x}}, \tilde{\tau}, \lambda^{\rm dam}) = \tilde{\pi}^\star(\lambda^{\rm dam})$ that selects the power purchased from the market as function of the market prices; thus, by definition, $\dot{Q}^{\rm b}(\lambda^{\rm dam})$ is directly defined by $\tilde{\pi}^\star(\lambda^{\rm dam})$.

B. RL for day-ahead and imbalance trading

As with MPC, the RL-based approach to trade in both markets consists of two separate RL algorithms:

- A first RL algorithm that trades with the day-ahead market. This is the algorithm proposed in Section V-A and it is agnostic of what happens in the imbalance market.
- A second RL algorithm that trades in the imbalance market and that considers the interaction with the day-ahead market. This algorithm runs in real time and it does not build bidding functions.

B.1 Training multiple RL agents

As each electricity market has its own rules and working principles, it is clear that a different RL agent for each market is needed. As an example, an RL agent for the imbalance market has a different state $\mathbf{s}$ as it knows more information than the agent for the day-ahead market, e.g. it knows the prices and allocations of the day-ahead market.

Based on this premise, when using RL to trade in two electricity markets, the problem becomes a multi-agent RL problem [51]. More specifically, as both agents are trying to minimize the economic cost, it becomes a collaborative multi-agent RL problem [52], [53].

While the literature has several methods for collaborative RL, e.g. join-action learners [52], we argue that the available methods might not be very suitable for the case of STESS. In particular, when training several agents at the same time, the environment becomes non-stationary [51], i.e. as each agent improves and changes its own policy the environment that the other agents perceive changes as well. This non-stationary condition invalidates the convergence properties of most single-agent RL algorithms [51]. While there are methods that address this by allowing every agent to observe the state and actions of the other agents, these are not applicable to STESSs. In particular, due to the sequential decision making nature of electricity markets, while the imbalance agent can know the state of the day-ahead agent, the opposite is not true, i.e. the information of the imbalance market is unknown at the time bids need to be submitted to the day-ahead market.

Based on the previous argument, we propose an RL approach for trading in the two markets where agents are not trained simultaneously. Instead, the day-ahead agent is trained first using the algorithm proposed in Section V-A, and the imbalance agent is trained afterwards including in its state information from the day-ahead market. This scheme has two benefits:

- **Convergence**: as the two RL agents are independently trained in two stationary environments, standard RL algorithms have guarantees of convergence.
- **Flexibility**: as the imbalance market is highly volatile, STESSs owners could potentially want to stop the trading in the imbalance market during periods of high volatility. As the agent for the day-ahead market is independent, STESSs could simply use the controls of this agent and be optimal in the more stable day-ahead market.

B.2 RL for the imbalance market

As the RL agent for the day-ahead is the same as the one described in Section V-A, we only need to define the RL agent that uses the information from the day-ahead and trades in the imbalance market. For the state space, besides the three values included in the state of the day-ahead agent, the new state includes past imbalance prices, past imbalance volumes, and the day-ahead price and energy allocation. In detail, at step k:

$$\mathbf{s}_k = (\mathbf{x}_k, \tau_k, \lambda_k^{\text{dam}}, \dot{Q}_k^{\text{dam}}, \lambda_{k-1}^{\text{imb}}, V_{k-1}^{\text{imb}}, \ldots, \lambda_{k-n_{\text{hrl}}}^{\text{imb}}, V_{k-n_{\text{hrl}}}^{\text{imb}}) \tag{13}$$

where V^{imb} represents the overall grid imbalance, and where the number of historical past values n_{hrl} is defined by the last lag uncorrelated to the imbalance price λ_k^{imb}. As an example, for The Netherlands, we observed $n_{\text{hrl}} = 3$ to be a good choice.

To define the action space $\mathbb{U}$, a single action $u \in \mathbb{R}^{n_{\text{in}}}$ has a similar format as before:

$$\mathbf{u} = (u_1, u_2, \ldots, u_{n_{\text{in}}}). \tag{14}$$

In detail, we consider that each input control u_i can take $n_{\text{dis}} + 1$ discrete values uniformly separated between -1 and 1. In particular, defining by $\dot{Q}_i^{\text{in,dam}}$ the energy purchased for storage device i in the day-ahead market, a value of $u_i = -1$ represents selling all the energy $\dot{Q}_i^{\text{in,dam}}$ in the imbalance market, i.e. $\dot{Q}_i^{\text{in}} = 0$. By contrast, a value of $u_i = 1$ represents buying all the energy that is still possible, i.e. $\dot{Q}_i^{\text{max}} - \dot{Q}_i^{\text{in,dam}}$, for storage device i, i.e. $\dot{Q}_i^{\text{in}} = \dot{Q}_i^{\text{max}}$. The selection of the output power is not considered as it is already selected by the day-ahead agent.

Besides the reward r that includes now the cost/income obtained in the imbalance market, and the simulation environment that also generates imbalance prices, the other parts of the RL agent remain the same.

B.3 Market interaction

In terms of the interaction with the agent for the day-ahead market, the STESS is controlled with both agents acting sequentially. First, one day-ahead, the day-ahead agent builds the bidding functions for the next day's day-ahead market. Next, the day-ahead market is cleared and the energy is allocated. Then, in real time, the imbalance agent uses the existing information of the day-ahead and imbalance markets to select the optimal power to buy/sell.

Unlike the agent for the day-ahead market, the imbalance agent does not build bidding functions as the imbalance market requires direct selection of the power $\dot{Q}^{\text{imb}}$ to buy/sell. As a result, the optimal policy $\pi^\star(\mathbf{s}_k)$ at time k directly selects the power to be traded based on available data $\mathbf{s}_k$ at that time step k, but not on the imbalance market price λ_k^{imb}.

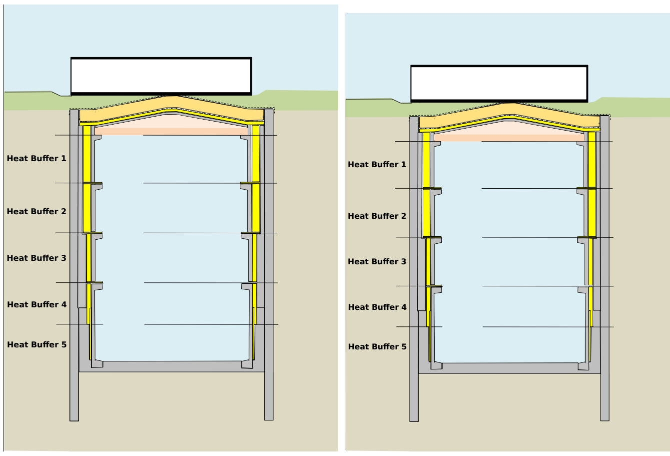

Fig. 1. Schematic representation of the STESS. **Left**: technical scheme representing the five heat buffers in the real system. **Right**: scheme representing the underground installation of the STESS.

VI. Case Study

To study the quality of the proposed control strategies, and in order to analyze the merits and disadvantages of each one of them, we consider the Ecovat vessel [54], a real SSTES. The system will be evaluated in eight case studies: first, the STESS will need to satisfy an uncertain heat demand during one year while minimizing the cost through the day-ahead market. Second, the STESS will need to supply the same heat demand but interacting with both the day-ahead and the imbalance market. For each of the two scenarios, we will consider two different heat demand profiles and two different countries.

A. Real STESS

The considered STESS is a large subterranean thermal stratified storage vessel with the ability to store heat for seasonal periods and to supply heat demand to a cluster of buildings. The system is divided into different segments or *heat buffers* that can be charged and discharged separately; the system has 5 thermal buffers with the top 4 buffers (see Figure 1) being able to be charged and discharged independently. Figure 1 provides a schematic representation of the vessel and Figure 2 illustrates the real system when it was under construction. For further details on the system we refer to [47].

B. System dynamics

The state of the STESS at time step k is defined by $\mathbf{x}_k = (T_{1,k}, T_{2,k}, T_{3,k}, T_{4,k}, T_{5,k})$, i.e. by the temperature stored in each of the 5 buffers as it is proportional to the stored energy. Similarly, as the top 4 buffers can be charged and discharged independently, the input and output power are respectively defined by $\dot{\mathbf{Q}}_k^{\text{in}} = (\dot{Q}_{1,k}^{\text{in}}, \ldots, \dot{Q}_{4,k}^{\text{in}})$ and

Fig. 2. Construction of the STESS. **Left**: installation of the last heat buffer. **Right**: STESS almost completely sealed.

$\dot{\mathbf{Q}}_k^{\text{out}} = (\dot{Q}_{1,k}^{\text{out}}, \ldots, \dot{Q}_{4,k}^{\text{out}})$. Finally, using the dynamical model for thermal stratified vessels proposed in [47], the dynamics of each heat buffer i at time k can be defined by:

$$T_{i,k+1} = T_{i,k} + a_1 \left(T_{i+1,k} + T_{i-1,k} - 2\,T_{i,k}\right) + a_2 \left(T_\infty - T_{i,k}\right) + a_3 (\dot{Q}_{i,k}^{\text{in}} - \dot{Q}_{i,k}^{\text{out}}), \quad (15)$$

where T_∞ represents the ambient temperature and is the only disturbance $\mathbf{d}$. For further details on the model we refer to [47]. Note that this is the dynamical model used for the RL simulator and for defining the dynamics constraint in the MPC.

C. Data

To set up the study, we consider the day-ahead and imbalance prices between 2015–2017 in The Netherlands[3], and the heat demand of a cluster of 5 buildings with a yearly-average heat demand of 220 MWh during the same time period[4]. As a second case study, we consider the day-ahead and the imbalance markets in Belgium, and the same heat demand.

The data of 2015 and 2016 is used as training data for the RL agents, and as the historical data for generating scenarios. The data of 2017 is used as out-of-sample data to evaluate the performance of the different algorithms.

D. Experimental Setup

To compare and study the control approaches, we evaluate their performance in terms of the economic cost that they incur when controlling the STESS for the full 2017 year in both The Netherlands and Belgium. As a baseline, we consider the economic cost of directly buying the instantaneous heat demand $\dot{Q}^{\text{d}}$ at the day-ahead market price. This baseline serves us to establish whether a control approach learns to trade energy, i.e. to study whether a control approach can use the STESS to reduce the energy cost. Moreover, to compare the algorithm in different conditions, the demand data is multiplied by 2 and used to evaluate the algorithms in the case of having 10 buildings, i.e. a yearly-average demand of 440 MWh.

The MPC algorithm is modeled using `Casadi` [55] and `python`, and then solved using *Ipopt* [56]. For the RL approach, the fitted-Q-iteration algorithm is implemented in `python` using the `Xgboost` [50] library. The forecaster of imbalance prices is also done via the `Xgboost` library.

It is important to note that, although both methods are based on completely different concepts, i.e. RL largely depends on the training data while MPC on the underlying optimization problem, the comparison between the methods is fair as the available data and dynamical model for both methods is exactly the same. In particular, MPC uses historical data to build price forecasts and RL uses the same historical data to build the simulation framework. Moreover, both methods consider the same dynamical model: MPC does it explicitly in the optimization problems while RL uses it in the simulation framework. While their solvers are different, this is the standard scenario in any comparison as different approaches have tailored solvers to the specific optimization problem, e.g. when comparing convex and non-convex models the convex models are estimated using a convex solver even though the non-convex models cannot make use of it.

[3]Collected from https://transparency.entsoe.eu/.

[4]Obtained from one of our research partners.

E. MPC approaches

To use the MPC approaches proposed in Section IV, a discrete set of prices has to be defined to build the bidding functions. To do so, we selected 15 discrete prices equally spaced between 0 and 70 €/MWh. This selection was done based on the price distribution in 2015–2016 and considering the computation time of solving a single OCP; however, a coarser or finer discretization could be used to respectively decrease the computation time or to increase the accuracy of the bidding functions. For prices above 70 €/MWh the bidding function was set to 0 considering the seldom occurrence of prices above this threshold. For negative prices, the bidding function was defined as the solution at 0 €/MWh.

The OCPs are defined by (8), (9), and (10), where:

- The dynamical constraint is represented by (15).
- The maximum power $\dot{\mathbf{Q}}_{\text{max}}^{\text{in}}$ to be traded in the market is defined by the electrical installation to charge the STESS. In our case $\dot{\mathbf{Q}}_{\text{max}}^{\text{in}} = 300\,\text{MW}$.
- The individual upper limits of charging and discharging, i.e. $g_{\text{in}}(\mathbf{x}_k)$ and $g_{\text{out}}(\mathbf{x}_k)$, are defined by the maximum heat transfer of the heat exchangers, which in turn is proportional to the temperature difference between the tank temperature and the temperature of the fluid in the heat exchangers.
- The limits on the STESS state are given by $\mathbf{x}_{\text{max}} = 286\,\text{K}$ and $\mathbf{x}_{\text{min}} = 263\,\text{K}$, where the lower limit is defined by the outer soil temperature and the upper limit by the safety margin to prevent water boiling in the tank.

F. RL approaches

The RL control algorithms proposed in Section V can be directly applied to the current case study:

- The time position τ is simply the day of the year.
- As the STESS has a seasonal cycle of a year, a RL episode length is defined as two years.
- The time-dependent constraints on the maximum power are implicitly enforced within the action space as the actions are normalized w.r.t. the maximum power.

G. Day-ahead market trading

The main results of the first study, i.e. the comparison of MPC and RL when only trading in the day-ahead market, are listed in Tables I and II. Table I displays the yearly economic cost when using both algorithms and the cost of not having an STESS, i.e. the cost of buying directly the heat demand in the day-ahead market; it also lists the economic savings of both algorithms w.r.t. the case of not having an STESS. Table II lists the offline costs, i.e. one-time computations, and online costs, i.e. real-time computations, of both algorithms.

TABLE I
MPC AND RL COMPARISON IN TERMS OF THEIR ECONOMIC COST WHEN ONLY TRADING IN THE DAY-AHEAD MARKET. THE SAVINGS ARE COMPUTED W.R.T. THE COST OF NOT HAVING AN STESS. FOR EACH CASE STUDY, THE BEST METHOD IS INDICATED IN BOLD.

		The Netherlands		Belgium	
		10 buildings	5 bldgs.	10 bldgs.	5 bldgs.
Cost [€]	No STESS	19384	9692	23490	11744
	MPC	**15206**	**6825**	**16826**	7033
	RL	15942	7465	17636	**7027**
Savings	MPC	**21.6%**	**29.6%**	**28.4%**	40.1%
	RL	17.8%	23.0%	24.9%	**40.2%**

TABLE II
MPC AND RL COMPARISON IN TERMS OF THEIR COMPUTATION TIME WHEN TRADING IN THE DAY-AHEAD MARKET. THE COMPARISON IS DONE IN TERMS OF ONLINE AND OFFLINE COMPUTATION TIME.

	Offline	Online
MPC	0	10–15 minutes
RL	1–2 days	<1 second

Independently of the country or heat demand level considered, the following observations can be made:

- Both algorithms can trade energy and make use of the STESS to reduce the economic cost. In particular, using the STESS and trading optimally, the algorithms can reduce the economic cost by 20–40%.
- The performance of both algorithms is similar, but MPC can obtain slightly lower costs and larger profits.
- While RL requires a long offline computation time, its cost online is almost negligible. In particular, as the optimal bidding functions are estimated offline, the computation time in real time is almost 0.
- By contrast, while MPC does not require offline computations, it needs 10–15 minutes in real time to build the bidding functions. However, as the bidding functions are submitted once per day and one day in advance, this large real-time computation cost does not represent a real problem/disadvantage.

Finally, to illustrate the generated bidding curves of both methods, Figure 3 displays the generated bidding curves the first day of the 5-buildings case study for the day-ahead market in The Netherlands. As it could be expected based on the results in Table I, both bidding curves are very similar.

H. Day-ahead and imbalance market trading

The main results of the second study, i.e. the comparison between MPC and RL when trading in both the day-ahead and

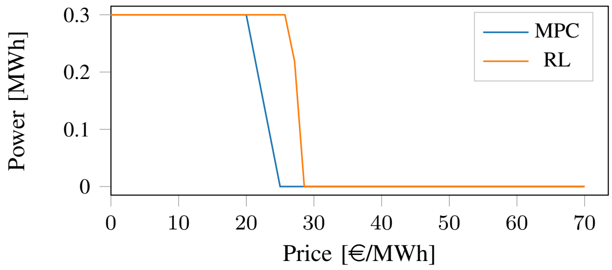

Fig. 3. Generated bidding curves by the MPC and RL algorithms on 01/01/2017 in The Netherlands when supplying heat for 5 buildings.

imbalance markets, are listed in Tables III, IV, and V. Table III displays the yearly economical cost and economic savings of both algorithms. Table IV lists their offline and online computation cost when trading in the imbalance market (the computation cost for trading in the day-ahead is the same as in Table II). As an extra comparison, Table V summarizes the percentage of times that each algorithm correctly up-regulates and down-regulates the grid, i.e. the percentage of times that the algorithm sells (buys) energy in the imbalance market while the TSO tries to up-regulate (down-regulate) the system.

TABLE III
MPC AND RL COMPARISON IN TERMS OF THEIR ECONOMIC COST WHEN TRADING IN THE DAY-AHEAD AND IMBALANCE MARKETS. THE SAVINGS ARE COMPUTED W.R.T. THE COST OF NOT HAVING AN STESS. FOR EACH CASE STUDY, THE BEST METHOD IS INDICATED IN BOLD.

		The Netherlands		Belgium	
		10 buildings	5 bldgs.	10 bldgs.	5 bldgs.
Cost [€]	No STESS	19384	9692	23490	11744
	MPC	**9227**	3544	**10569**	4401
	RL	11176	**3437**	11468	**3872**
Savings	MPC	**52.4%**	63.4%	**55.0%**	62.5%
	RL	42.3%	**64.5%**	51.8%	**67.0%**

TABLE IV
COMPUTATION COST OF THE MPC AND RL APPROACHES WHEN TRADING IN THE IMBALANCE MARKET. THE COMPARISON IS DONE IN TERMS OF ONLINE AND OFFLINE COMPUTATION TIME.

	Offline	Online
MPC	0	30–45 seconds
RL	1–2 days	<1 second

As before, independently of the case study considered, the following observations can be made:

- As for day-ahead trading, both algorithms perform very similar to each other. However, unlike in the case of only day-ahead trading, MPC no longer performs slightly better. Instead, RL performs slightly better for lower heat demand profiles (5 buildings), and MPC performs better for higher heat demand profiles (10 buildings).
- Trading in both markets is much more beneficial than trading only in the day-ahead market as the costs are

TABLE V
MPC AND RL COMPARISON IN TERMS OF THE % OF TIMES THAT THEY CORRECTLY UP-REGULATE OR DOWN-REGULATE THE GRID, I.E. % OF TIMES THAT THEY SELL/BUY ENERGY IN THE IMBALANCE MARKET WHEN THE TSO UP/DOWN-REGULATES. FOR EACH CASE STUDY, THE BEST METHOD IS INDICATED IN BOLD.

		The Netherlands		Belgium	
		10 buildings	5 bldgs.	10 bldgs.	5 bldgs.
Up-regul.	MPC	**51%**	47%	44%	47%
	RL	49%	46%	**50%**	**52%**
Down-reg.	MPC	70%	66%	68%	55%
	RL	**81%**	**81%**	**81%**	**80%**

halved w.r.t. day-ahead trading. In particular, while day-ahead trading reduces the economic cost by 20–40%, trading in the two market reduces the cost up to 60–70%.

- As before, RL requires large offline computation costs but negligible online computation costs. By contrast, MPC has no offline computation costs but requires 30–45 seconds to obtain the optimal trading strategy for the imbalance market. Since the imbalance market is cleared every 15 minutes and optimal decisions are made within seconds, it can be argued that the online computation cost of MPC might now represent a problem.
- When buying energy in the imbalance market, the RL algorithm helps the TSO to down-regulate the grid. In particular, approximately 80% of the times the RL algorithm buys energy, the TSO simultaneously tries to reduce the grid generation or to increase the grid consumption. While the MPC algorithm also helps, this contribution is worse as it only helps to down-regulate 55–70% of the time.
- By contrast, when selling energy in the imbalance market, none of the algorithms help much to up-regulate: only 45–55% of the times an algorithm sells energy the TSO is simultaneously trying to up-regulate.

VII. DISCUSSION

In this section, based on the obtained results, we discuss the merits and disadvantages of the proposed control approaches, the benefits of using STESSs for energy trading, how to optimally operate STESSs to maximize their profits, and the generality and optimality of the proposed methods.

A. Merits of each control approach

We start the discussion by analyzing the merits of the different proposed approaches in the two trading contexts.

A.1 Day-ahead trading

When trading only in the day-ahead market, both approaches can trade energy with a similar performance despite their underlying differences. Therefore, while MPC obtains slightly lower economic costs than RL, it is necessary to consider other metrics in order to make a meaningful comparison.

When considering the online computation time, both algorithms are feasible for real-life applications. Thus, the largest difference between both approaches is the offline computation time. While this metric does not play a role most of the time, i.e. it usually represents one-time computation costs, it might be important when the system regularly goes under maintenance, something breaks down, or the market has a big change. In particular, if any of these events happens, MPC can easily adapt itself by a change in the OCP or by re-estimating the dynamical model (which does not take more than some minutes). By contrast, RL requires 1–2 days to re-estimate the optimal policy under the new conditions, which hinders the day-ahead trading. Thus, MPC has in general better adaptability to environmental conditions.

Based on this analysis, it becomes clear that MPC is a better approach when trading only in the day-ahead market. Particularly, slightly better optimal solutions together with a better adaptability to environmental changes make the proposed MPC approach a better solution in this case.

A.2 Day-ahead and imbalance trading

Similar to the case of only day-ahead trading, when trading in the day-ahead and imbalance market the two proposed approaches obtain good solutions. In particular, while RL performs slightly better for lower heat demand profiles (5 buildings) and MPC performs better for higher heat demand profiles (10 buildings), these difference are not very large and, as before, other metrics need to be considered.

While the online computation time for day-ahead trading was not an issue, for the case of imbalance trading it becomes one. In detail, due to the real-time nature of the imbalance market, optimal decisions should be made in seconds. As the proposed MPC approach requires 30–45 seconds to compute an optimal solution, it can potentially fail to provide an optimal trading strategy.

As a result, while the proposed MPC approach still has a better adaptability to environmental changes, one could argue that it is a less appropriate control strategy than the proposed RL approach. The latter, with its negligible real-time computation cost, equal quality solutions, and better regulatory capabilities, is a better choice when it comes to trading in the imbalance market.

B. The importance of market trading for STESSs

Based on the obtained results, it is clear that optimal control approaches, either MPC or RL, are key to maximize the profits of STESSs and to ensure their widespread use as optimal control strategies and can reduce the energy cost by 60–70%. In this context, the largest profits are obtained when the STESS trades in multiple markets. In particular, while a traditional STESS would restrict its trading to the day-ahead market to avoid unnecessary risks, in this paper we show that STESSs can dramatically reduce their costs by using optimal control strategies and trading also in the imbalance market.

C. STESSs as regulation tools

Looking at the results of Table V, it can be argued that the economic goal of STESSs is (partially) aligned with the regulatory duties of the TSO. In particular, in the case of RL, 80% of the times the STESS buys energy in the imbalance market, it helps the TSO to down-regulate the system. This behavior is seen for the various case studies considered, which included different imbalance markets and different heat demands. In the case of MPC, this effect is not so pronounced; nevertheless, it sill helps the TSO 55–70% of the times.

While the same cannot be said about up-regulation, i.e. only 50% of the times the STESS sells energy in the imbalance market it is actually helping up-regulate the grid, it can be argued that wrongly up-regulating is less critical than wrongly down-regulating. In particular, if the STESS wrongly sells energy in the imbalance market, the TSO can always request somebody to reduce their generation, i.e. down-regulate. However, if the STESS wrongly buys energy in the imbalance market, the TSO has to request somebody to increase their generation; as the generation is limited, there might not be an available agent that can provide that service.

As additional remark, to further improve the regulatory services of STESSs, communication between the TSO and the STESS could be established. In particular, in the current setup, the STESS simply optimizes its profit without considering the TSO. Thus, to improve this, the TSO could simply indicate the STESS whether it is allow to buy or sell energy, i.e. whether the TSO plans to down or up-regulate, and the STESS could take its optimal action if it helps the TSO and its own profit.

D. Generality of the methods

While the case study focused on a specific STESS, i.e. a latent heat storage via water stratification, the proposed methods are general and can be applied to any STESS. Indeed, with the proposed methods, the several challenges that prevent the development of efficient control solutions for STESS trading can be tackled, namely: scenario generation and quantification of price uncertainty for long horizons, small computation costs for real-time control, and adaptability to market changes.

E. Optimality of the methods

The optimality property of the proposed methods is affected by the following elements: 1) the optimization problems are non-convex, 2) the quality of the solutions depend on the accuracy of the forecasting method, and 3) in a multi-stage optimization problem the decision taken at the first stage will have an effect in future stages. In this context, it is important to remark that the methods are nonetheless optimal from the perspective that they take a local optimal solution at every state with the information that is known:

- The first optimization problem takes an optimal decision considering that at the moment of the decision only knows a forecast of the future prices.
- The second optimization problem takes an optimal decision with updated information and considering that market conditions have been changed. While this decision may differ from the first optimal solution, the solution is nonetheless a local optimum at the time that the decision is made.

Within the same context of optimality, to evaluate the proposed methods, the obtained solutions should ideally be compared with the real optimal solutions considering perfect knowledge of the future. However, this is not possible nor fair for two reasons:

1) The optimization problem that provides the optimal solution is non-convex. Therefore, such an analysis would involve comparing two local minima and it would not involve a real optimal baseline.
2) The proposed approaches need to rely on forecasting methods while the baseline solution have perfect knowledge of the future. In this context, the quality of the proposed methods depend on an external factor (forecasts) that the baseline solution does not.

VIII. Conclusions

We have proposed several optimal control strategies for *seasonal thermal storage systems (STESSs)* when interacting with electricity markets. Particularly, while in the literature there are control strategies for STESSs and there are optimal trading strategies for traditional storage systems, the former do not allow STESSs to trade in the markets and the latter are not suitable for STESSs. To fill that gap, we have proposed a *model predictive control (MPC)* and a *reinforcement learning (RL)* approach for the case of having an STESS trading in the day-ahead electricity market. In addition, we argued that trading in one market is not optimal, and proposed another MPC and another RL approach for the case of having an STESS trading in both the day-ahead and the imbalance markets.

Based on a case study involving a real STESS, it was shown that, despite the similarity in the optimal solutions of the proposed algorithms, MPC is a better trading strategy for the day-ahead market due to its larger adaptability. In contrast, for trading in the imbalance market, the proposed RL approach is a more suitable control strategy as it has negligible real-time computation costs, leads to similar economic costs as MPC, and has better regulatory capabilities.

It was also shown that STESS are potential tools for grid regulation and that the economic incentive of STESSs are aligned with the regulatory duties of TSOs. Similarly, it was demonstrated that optimal control strategies are needed to optimize the profit of STESSs and to ensure their widespread use.

In future research, we intend to further explore the use of STESSs as regulation devices. Moreover, as stochastic approaches can further improve the performance of the control algorithms in the context of long horizons, we will analyze the advantages of using stochastic MPC approaches for seasonal storage systems. Finally, we will also study the trade-offs between MPC and RL to obtain a set of generalizable trade-offs that are independent from the case study.

Acknowledgments

The authors would also like to thank Dr. Fjo De Ridder for the fruitful discussions, and the people at Ecovat for granting the access to their seasonal thermal energy storage system.

LIST OF SYMBOLS

Type	Symbol	Definition
Indices	t	Continuous time index
	k	Discrete time index
	h	Discrete hourly time index
Dynamic Systems	x	State of a general dynamic system
	u	Controls of a general dynamic system
	d	Disturbances of a general dynamic system
	$\dot{Q}^{\mathrm{in}}$	Input power
	$\dot{Q}^{\mathrm{out}}$	Output power
	$\dot{Q}^{\mathrm{d}}$	Heat demand
	$\dot{Q}^{\mathrm{max}}$	Maximum input power
	n_{in}	Number of inputs
	n_{out}	Number of outputs
	n_{units}	Number of individual storage units
Electricity Markets	$\dot{Q}^{\mathrm{m}}$	Allocated power from a general market
	$\dot{Q}^{\mathrm{dam}}$	Allocated power from the day-ahead market
	$\dot{Q}_i^{\mathrm{in,dam}}$	Power for device i from the day-ahead market
	$\dot{Q}^{\mathrm{imb}}$	Allocated power from the imbalance market
	$\dot{Q}^{\mathrm{b}}(\cdot)$	Biding function for electricity market
	λ	General price
	λ^{dam}	Price in the day-ahead market
	λ^{imb}	Price in the imbalance market
	V^{imb}	Volume of the imbalance
Forecast	$\bar{\dot{Q}}^{\mathrm{d}}$	Generated scenarios of heat demand
	$\bar{\lambda}^{\mathrm{dam}}$	Generated scenarios of day-ahead prices
	$\bar{\lambda}^{\mathrm{imb}}$	Generated scenarios of imbalance prices
	$\hat{\lambda}^{\mathrm{imb}}$	Forecast of imbalance market prices
MPC	N	Number of discrete time intervals
	T	Optimization horizon
	λ^{b}	Discrete price for building bidding functions
	n_{p}	Number of discrete price in bidding functions
RL	s	State of agent
	u	Action taken by agent
	$\mathbb{U}$	Discrete set of possible actions
	r	Reward obtained when taking action
	π	Agent policy used to take actions
	T_{e}	Episode length
	τ	Seasonal time index in the agent state
	n_{hrl}	Number of past lags in the RL state
	n_{dis}	Number of discretized inputs

APPENDIX A
SCENARIO GENERATION

When defining the MPC algorithms, it was assumed that the expected day-ahead prices $\{\bar{\lambda}_k^{\mathrm{dam}}\}_{k=1}^N$, imbalance prices $\{\bar{\lambda}_k^{\mathrm{imb}}\}_{k=1}^N$, heat demand values $\{\bar{\dot{Q}}_k^{\mathrm{d}}\}_{k=1}^N$, and disturbances $\{\bar{\mathbf{d}}_k\}_{k=1}^N$ were given. In this appendix, the methodology to generate these time series is explained.

A. Motivation

In order for the MPC to provide good solutions, the expected time series have to be realistic. Therefore, any forecasting method for these time series has to model the time correlation of a single time series and the inter-correlation between the different times series. While there are methods in the literature to create those forecasts, these are limited to short-term horizons with small resolutions, e.g. hourly, or long-term horizons with broad resolutions, e.g. daily, [57]. The main problem of generating forecasts with small resolutions and long horizons is the accuracy: due to the large uncertainty, it is nearly impossible to forecast electricity prices or loads with an hourly resolution one year in advance. Instead of forecasting the expected value, one could generate a set of different scenarios representing possible future realizations. However, as with the literature of forecasts, the field of scenario generation has, to the best of our knowledge, no reliable method to generate long-term scenarios with small resolutions. In particular, the literature of scenario generation for correlated time series is limited to short-term horizons [44]–[46]. In this case, the problem is computational tractability: whether the methods are based on trees [58] or on copulas [44]–[46], the computational cost is too large, e.g. in the case of trees the number of scenarios grows exponentially with the horizon [59].

B. Method

In this paper, we propose a very simple, yet useful, method to generate scenarios of correlated time series for long-term horizons. Then, we use the average of the scenarios at every time point as the expected values used in the MPC. The only two requirements of the proposed method are: 1) to have at least as many historical data as the horizon length of the scenarios; 2) to have historical data with a resolution equal to or lower than the resolution of the scenarios.

Given a set of n_{ts} historical time series of length n_{h}, i.e. $\{\mathbf{x}_{1,j}\}_{j=1}^{n_{\mathrm{h}}}$, $\{\mathbf{x}_{2,j}\}_{j=1}^{n_{\mathrm{h}}}, \ldots, \{\mathbf{x}_{n_{\mathrm{ts}},j}\}_{j=1}^{n_{\mathrm{h}}}$, the proposed method generates any number n_{s} of future scenarios of length $N \leq n_{\mathrm{h}}$:

$$\left\{\{\bar{\mathbf{x}}_{1,k}^i\}_{k=n_{\mathrm{h}}+1}^{n_{\mathrm{h}}+N}, \{\bar{\mathbf{x}}_{2,k}^i\}_{k=n_{\mathrm{h}}+1}^{n_{\mathrm{h}}+N}, \ldots, \{\mathbf{x}_{n_{\mathrm{ts}},k}^i\}_{k=n_{\mathrm{h}}+1}^{n_{\mathrm{h}}+N}\right\}_{i=1}^{n_{\mathrm{s}}}, \tag{16}$$

In detail, the method consists of 7 steps:

1) Select a representative horizon $N' << N$ so that, given any two time series of length N', any point after N' is uncorrelated with the first point of both time series. In the case of day-ahead market prices, hourly resolution, and a year horizon, we empirically observed $N' = 8$ days to be a good choice as N' includes the weekly and daily seasonalities correlations w.r.t. the first price.
2) Define a subset $\{j_1, j_2, \ldots\}$ of past indices whose associated values are correlated to the expected values at time step $k = 1$. For time series with seasonalities, these indices represent past values at lags equal to multiples of these seasons, e.g. for time series with daily and weekly seasonalities these indices could represent the values 1 day and 1 week in the past. In general, one could use a correlation study to determine the relevant indices.
3) Sample from the historical dataset a subset $\{\mathbf{x}_{1,j}\}_{j=j_i}^{j_i+N'}$, $\{\mathbf{x}_{2,j}\}_{j=j_i}^{j_i+N'}, \ldots, \{\mathbf{x}_{n_{\mathrm{ts}},j}\}_{j=j_i}^{j_i+N'}$, where the time index j_i is randomly selected from the past indices $\{j_1, j_2, \ldots\}$.
4) Use the previous sample as the first N' points $\{\bar{\mathbf{x}}_{1,k}^1\}_{k=1}^{N'}$, $\{\bar{\mathbf{x}}_{2,k}^1\}_{k=1}^{N'}, \ldots, \{\mathbf{x}_{n_{\mathrm{ts}},k}^1\}_{k=1}^{N'}$ of the first scenario.

5) Repeat steps 2–4 but defining the subset $\{j_1, j_2, \ldots\}$ as the time indices correlated to the next time point in the scenario, i.e. $k = N' + 1$.
6) Repeat step 5 until a whole scenario is obtained, i.e. N/N' times.
7) Repeat 6 until the n_s scenarios are obtained.

It is important to note that, depending on the application, the selection of j_i will vary. For example, in the case of electricity prices and $N' = 1\,$week, considering that prices have yearly, weekly and daily seasonalities, we observed that a good choice for the subset $\{j_1, j_2, \ldots\}$ are time indices representing 1, 2, 50, 51, 52, 53, 54, 102, 103, 104, 105, or 106 weeks in the past, i.e. past time indices that respect the yearly and weekly seasonality (the indices represent the last 2 weeks, last year $\pm$ 2 weeks, and 2 years ago $\pm$ 2 weeks). In a more general setup, one could use a correlation study or a method like k-nearest neighbors [60] to determine the relevant past indices.

Appendix B
Imbalance price forecast

As explained in Section IV, a forecast of the imbalance price $\hat{\lambda}_1^{\mathrm{imb}}$ at the first time step of the MPC algorithm is needed. In particular, as the MPC algorithm for the imbalance market decides the traded power directly based on $\hat{\lambda}_1^{\mathrm{imb}}$, it is important for $\hat{\lambda}_1^{\mathrm{imb}}$ to be as accurate as possible. While the expected imbalance price $\bar{\lambda}_1^{\mathrm{imb}}$ obtained from the scenario generation method could be used as a forecast, this is not the most accurate prediction as the scenario generation method simply resamples from past data and does necessarily consider the most recent information.

As the literature of electricity price forecasting does not contain, to the best of our knowledge, a method for imbalance price forecasting, in this paper we propose a first method for it. In detail, the boosting trees model [50] is selected as the forecasting model due to its simplicity and recent success in forecasting day-ahead prices [3]. As input features, the model considers:

- The last n_1 imbalance prices, where n_1 is optimized.
- The last n_2 imbalance volumes, where n_2 is optimized.
- The day-ahead electricity price at the hour of interest.
- The hour of the day and the day of the week.

The hyperparameters of the boosting tree model, e.g. number of trees, are simultaneously optimized with n_1 and n_2 using the tree-Parzen estimator [61]. The selection of this algorithm to do the feature and hyperparameter selection is motivated by its recent success in other energy applications [3], [39]. In this context, it is important to note that, as any method for forecasting time series data, no performances guarantees can be provided as the accuracy of any forecasting method will always depend upon the data under study. Moreover, since the main goal of the paper is the control solutions and not the forecasting methods, a comparison between the proposed method and typical time series methods, e.g. ARIMA, is not presented. However, for a thorough comparison of different forecasting methods for electricity prices we refer the interested reader to [3].

As a final remark, it is important to note that, while this is the first method for forecasting imbalance prices, there exist other methods to forecast real-time *local marginal prices (LMPs)* [40], [41], [62]. However, real-time LMPs have different characteristics and represent a different concept than imbalance prices. In particular, the volatility of imbalance prices is larger than real-time LMPs; thus, forecasting imbalance prices is arguably harder than forecasting real-time LMPs (in our experience, forecasting imbalance prices with a horizon larger than one hour is nearly impossible; however, methods for real-time LMPs usually have forecasting horizons up to 6 hours). To demonstrate this argument, we will compare the difference in volatility during the whole year 2019 between the real-time LMPs in the *Pennsylvania-Jersey-Maryland (PJM)* market and the imbalance prices in the Dutch market. Then, we will analyze two time series related to volatility:

1) The spread between the day-ahead and the real-time/imbalance market: the lower the spread, the more stable the real-time/imbalance prices are.
2) The price difference $\Delta\lambda_k$ between successive prices λ_{k+1} and λ_k: low volatility is characterized by low variations between successive prices.

Table VI represents four different statistics for the two time series. Since volatility is determined by the absolute size of the spread and $\Delta\lambda_k$ (but not by their sign), the mean is computed as the mean of the absolute value. Three main observations can be made:

- The difference between successive prices is much larger for the imbalance price. This illustrates that the volatility of imbalance prices is larger, and thus, that imbalance prices are harder to forecast.
- The spread and its variation are much larger for the imbalance prices. Thus, this also shows that imbalance prices are harder to forecast.
- The maximum deviations (both for spread and $\Delta\lambda_k$) are larger for imbalance prices. This means that imbalance prices have larger price spikes, which in turn means that forecasting imbalance prices is harder.

TABLE VI
Comparing the volatility between the PJM real-time market and the Dutch imbalance market in 2019. The mean is computed as the mean of the absolute value.

	Spread		$\mathbf{\Delta\lambda_k}$	
	Real-time	**Imbalance**	**Real-time**	**Imbalance**
Mean	4.79	26.12	9.05	24.94
Std	18.69	44.31	20.83	45.81
Max	81.38	525.65	70.56	526.5
Min	-630.69	-832.42	-672.26	-897.27

Appendix C
Electricity markets

In this appendix, we introduce the electricity markets considered in this paper. In particular, while electricity is traded in different markets, in this paper we focus on the day-ahead and imbalance markets.

A. Trading in the day-ahead market

The day-ahead electricity market [63] is a type of power exchange widely used around the world. In this market, consumers/producers submit bids for day d before some deadline on day $d-1$, where a bid indicates how much they are willing to pay/ask for different power volumes. With some exceptions, these bids are usually hourly based, i.e. each market player submits 24 bids. After the deadline, the market operator takes into account all the bids and computes the market clearing price for each of the 24 hours. Then, consumer/producer bids higher/lower than or equal to the market clearing prices are approved, and a contract is established.

B. Trading in the imbalance market

Electricity, unlike most commodities, cannot be stored in very large amounts and requires almost equal consumption and generation at all times [64], [65]. However, as electricity consumption and generation are uncertain and weather dependent, in practice there are always imbalances between generation and consumption created by market agents that do not consume or generate what they had promised in the markets [66]. These imbalances have an adverse impact on the electrical grid frequency [66], and can lead to grid problems and instabilities and in some cases even blackouts [64]. To correct these imbalances, the *transmission system operator (TSO)* manages a so-called reserve market [65], [67]. On this market, specific market agents sell their available energy reserves to the TSO, i.e. their capacity to reduce and increase their generation and consumption, and the TSO purchases some of these reserves days or weeks in advance. Then, the TSO activates the required reserves on real time to correct grid imbalances.

Based on the price the TSO had to pay to correct the imbalances, it invoices the market agent that have caused the imbalance [68]. This mechanism where all the market agents pay for their imbalances is known as imbalance market or imbalance settlement [68]. Usually, this market is cleared every 15 minutes, i.e. each market agent pays for their cumulative positive or negative imbalance in intervals of 15 minutes.

In some countries, it is discouraged or even forbidden to use such a market for electricity trade, i.e. market agents are expected to trade honestly in the markets available before delivering time and only produce unexpected imbalances. However, as during periods of positive imbalances, i.e. when generation is larger than consumption, prices are low, and during period of negative imbalances prices are high, the economic incentive of market agents in this imbalance market is aligned with the regulatory duties of the TSO. Based on that, some other countries, e.g. The Netherlands [69], allow and encourage participation in this market.

APPENDIX D
CONTROL ALGORITHMS

In this appendix, we introduce the control algorithms considered in the paper. Particularly, we focus on the two most important state-of-the-art control families: predictive control via MPC and artificial intelligence via RL.

A. Introduction to MPC

The general idea of MPC is to, at each discrete time step k, obtain the optimal control $\mathbf{u}_k^\star$ by using the following iterative structure:

1) Read current state $\mathbf{x}_k$.
2) Based on $\mathbf{x}_k$, solve the relevant *optimal control problem (OCP)* over a horizon of N time intervals.
3) Based on the solution of this optimization problem, obtain the optimal control $\mathbf{u}_k^\star$.
4) Apply this control to the system.
5) Repeat the process again for the next time step $k+1$.

For more details on MPC we refer to [42].

B. Introduction to RL

The general idea of RL is to, at each time step k, obtain the optimal control $\mathbf{u}_k^\star$ by using an optimal policy $\pi^\star(\mathbf{s}_k)$ (a function that outputs the optimal action $\mathbf{u}_k^\star$ for each state $\mathbf{s}_k$). To learn the policy $\pi^\star(\mathbf{s}_k)$, the RL algorithm assumes that the dynamical system and its environment can be modeled via a Markov decision process [43], [49]. In detail, it assumes that:

- The system lives in a discrete-time world.
- The system is controlled by an agent that takes actions $\mathbf{u}$ among a discrete set of actions $\mathbb{U} = \{\mathbf{u}^1, \ldots \mathbf{u}^{n_\text{a}}\}$.
- The system and the environment are modeled by the agent state $\mathbf{s}$ where, in general, the state $\mathbf{x}$ of the system is part of the state $\mathbf{s}$ of the RL agent.
- At every discrete time step k, the agent takes an action $\mathbf{u}_k$ and transitions from state $\mathbf{s}_k$ to $\mathbf{s}_{k+1}$ based on some probabilistic dynamics $p(\mathbf{s}_{k+1}|\mathbf{s}_k, \mathbf{u}_k)$.
- In the transition, the agent receives a reward r_k based on a distribution $q(r_k|\mathbf{s}_k, \mathbf{u}_k)$.

During training, the RL agent iteratively performs an exploration step and an exploitation step:

- **Exploration**: the agent controls the system to interact with the environment for a number of n_steps steps. Then, it gathers the data resulting from that interaction in a memory dataset $\mathbb{M} = \{(\mathbf{s}_k, \mathbf{u}_k, \mathbf{s}_{k+1}, r_k\}_{k=1}^{n_\text{steps}}$. To select $\mathbf{u}_k$ during the exploration, the agent uses both optimal actions from $\pi^\star(\mathbf{s}_k)$ and random actions. In future repetitions of this exploration step, new data is added to the memory $\mathbb{M}$.
- **Exploitation**: the agent uses $\mathbb{M}$ to improve the optimal policy $\pi^\star(\mathbf{s}_k)$. In particular, $\pi^\star(\mathbf{s}_k)$ is estimated so that the expected value of the cumulative sum of discounted rewards R is maximized:

$$R = \sum_{k=1}^{T_\text{e}} \gamma^{T_\text{e}-k} \mathop{\mathbb{E}}_{q(r_k|\mathbf{s}_k,\mathbf{u}_k)} \{r_k\}, \tag{17}$$

  where T_e is the length of a RL episode, i.e. for how long the RL agent takes decisions, and γ is a discount factor that prioritizes earlier rewards and allows R to be finite even for episodes with an infinite horizon.

In general, what defines and separates the large family of RL algorithms is the manner in which these two steps are performed, i.e. the number of steps n_steps, the size of the memory $\mathbb{M}$, or the algorithm to estimate $\pi^\star(\mathbf{s}_k)$.

REFERENCES

[1] B. K. Sovacool, "How long will it take? Conceptualizing the temporal dynamics of energy transitions," *Energy Research & Social Science*, vol. 13, pp. 202–215, 2016.

[2] U. Bardi, "The grand challenge of the energy transition," *Frontiers in Energy Research*, vol. 1, p. 2, 2013.

[3] J. Lago, F. De Ridder, and B. De Schutter, "Forecasting spot electricity prices: deep learning approaches and empirical comparison of traditional algorithms," *Applied Energy*, vol. 221, pp. 386–405, 2018.

[4] IRENA, "Electricity storage and renewables: Costs and markets to 2030," International Renewable Energy Agency, Tech. Rep., 2017.

[5] S. Ugarte, J. Larkin, B. van der Ree, V. Swinkels, M. Voogt, N. Friedrichsen, J. Michaelis, A. Thielmann, M. Wietschel, and R. Villafáfila, "Energy storage: Which market designs and regulatory incentives are needed?" European Parliament's Committee on Industry, Research and Energy (ITRE), Tech. Rep., 2015.

[6] O. Schmidt, S. Melchior, A. Hawkes, and I. Staffell, "Projecting the future levelized cost of electricity storage technologies," *Joule*, vol. 3, no. 1, pp. 81–100, 2019.

[7] Agora Verkehrswende, "The future cost of electricity-based synthetic fuels," Agora Energiewende and Frontier Economics, Tech. Rep., 2018.

[8] A. Lucas and S. Chondrogiannis, "Smart grid energy storage controller for frequency regulation and peak shaving, using a vanadium redox flow battery," *International Journal of Electrical Power & Energy Systems*, vol. 80, pp. 26–36, 2016.

[9] M. Walter, M. V. Kovalenko, and K. V. Kravchyk, "Challenges and benefits of post-lithium-ion batteries," *New Journal of Chemistry*, 2020.

[10] J. W. Choi and D. Aurbach, "Promise and reality of post-lithium-ion batteries with high energy densities," *Nature Reviews Materials*, vol. 1, no. 4, 2016.

[11] J. O. Abe, A. P. I. Popoola, E. Ajenifuja, and O. Popoola, "Hydrogen energy, economy and storage: Review and recommendation," *International Journal of Hydrogen Energy*, vol. 44, pp. 15 072–15 086, 2019.

[12] O. Pesonen and T. Alakunnas, "Energy storage a missing piece of the puzzle for the self-sufficient living," Lapland Unniversity of Applied Sciences, Tech. Rep., 2017.

[13] B. Battke and T. S. Schmidt, "Cost-efficient demand-pull policies for multi-purpose technologies the case of stationary electricity storage," *SSRN Electronic Journal*, 2015.

[14] J. Xu, R. Wang, and Y. Li, "A review of available technologies for seasonal thermal energy storage," *Solar Energy*, vol. 103, pp. 610–638, 2014.

[15] I. Sarbu and C. Sebarchievici, "A comprehensive review of thermal energy storage," *Sustainability*, vol. 10, no. 1, p. 191, 2018.

[16] G. Darivianakis, A. Eichler, R. S. Smith, and J. Lygeros, "A data-driven stochastic optimization approach to the seasonal storage energy management," *IEEE control systems letters*, vol. 1, no. 2, pp. 394–399, 2017.

[17] V. Rostampour and T. Keviczky, "Probabilistic energy management for building climate comfort in smart thermal grids with seasonal storage systems," *IEEE Transactions on Smart Grid*, vol. 10, no. 4, pp. 3687–3697, 2019.

[18] E. Saloux and J. A. Candanedo, "Control-oriented model of a solar community with seasonal thermal energy storage: development, calibration and validation," *Journal of Building Performance Simulation*, pp. 1–23, 2018.

[19] J. Arteaga and H. Zareipour, "A price-maker/price-taker model for the operation of battery storage systems in electricity markets," *IEEE Transactions on Smart Grid*, vol. 10, no. 6, pp. 6912–6920, 2019.

[20] N. Yu and B. Foggo, "Stochastic valuation of energy storage in wholesale power markets," *Energy Economics*, vol. 64, pp. 177–185, 2017.

[21] M. Kazemi, H. Zareipour, N. Amjady, W. D. Rosehart, and M. Ehsan, "Operation scheduling of battery storage systems in joint energy and ancillary services markets," *IEEE Transactions on Sustainable Energy*, vol. 8, no. 4, pp. 1726–1735, 2017.

[22] H. Khani, R. K. Varma, M. R. D. Zadeh, and A. H. Hajimiragha, "A real-time multistep optimization-based model for scheduling of storage-based large-scale electricity consumers in a wholesale market," *IEEE Transactions on Sustainable Energy*, vol. 8, no. 2, pp. 836–845, 2017.

[23] A. Gonzalez-Garrido, A. S. de Ibarra, H. Gaztanaga, A. Milo, and P. Eguia, "Annual optimized bidding and operation strategy in energy and secondary reserve markets for solar plants with storage systems," *IEEE Transactions on Power Systems*, vol. 34, no. 6, pp. 5115–5124, 2019.

[24] S. Nojavan, A. Akbari-Dibavar, and K. Zare, "Optimal energy management of compressed air energy storage in day-ahead and real-time energy markets," *IET Generation, Transmission & Distribution*, vol. 13, no. 16, pp. 3673–3679, 2019.

[25] A. Akbari-Dibavar, K. Zare, and S. Nojavan, "A hybrid stochastic-robust optimization approach for energy storage arbitrage in day-ahead and real-time markets," *Sustainable Cities and Society*, vol. 49, p. 101600, 2019.

[26] X. Fang, B.-M. Hodge, L. Bai, H. Cui, and F. Li, "Mean-variance optimization-based energy storage scheduling considering day-ahead and real-time LMP uncertainties," *IEEE Transactions on Power Systems*, vol. 33, no. 6, pp. 7292–7295, 2018.

[27] E. Nasrolahpour, J. Kazempour, H. Zareipour, and W. D. Rosehart, "A bilevel model for participation of a storage system in energy and reserve markets," *IEEE Transactions on Sustainable Energy*, vol. 9, no. 2, pp. 582–598, 2018.

[28] S. Karhinen and H. Huuki, "Private and social benefits of a pumped hydro energy storage with increasing amount of wind power," *Energy Economics*, vol. 81, pp. 942–959, 2019.

[29] B. Cheng and W. Powell, "Co-optimizing battery storage for the frequency regulation and energy arbitrage using multi-scale dynamic programming," *IEEE Transactions on Smart Grid*, pp. 1–1, 2016.

[30] S. Grillo, A. Pievatolo, and E. Tironi, "Optimal storage scheduling using markov decision processes," *IEEE Transactions on Sustainable Energy*, vol. 7, no. 2, pp. 755–764, 2016.

[31] R. Bellman, *Dynamic Programming*. Princeton University Press, 1957.

[32] A. Oudalov, D. Chartouni, and C. Ohler, "Optimizing a battery energy storage system for primary frequency control," *IEEE Transactions on Power Systems*, vol. 22, no. 3, pp. 1259–1266, 2007.

[33] P. Zou, Q. Chen, Q. Xia, G. He, and C. Kang, "Evaluating the contribution of energy storages to support large-scale renewable generation in joint energy and ancillary service markets," *IEEE Transactions on Sustainable Energy*, vol. 7, no. 2, pp. 808–818, 2016.

[34] V. Rostampour, M. Jaxa-Rozen, M. Bloemendal, and T. Keviczky, "Building climate energy management in smart thermal grids via aquifer thermal energy storage systems," *Energy Procedia*, vol. 97, pp. 59–66, 2016.

[35] F. De Ridder, M. Diehl, G. Mulder, J. Desmedt, and J. Van Bael, "An optimal control algorithm for borehole thermal energy storage systems," *Energy and Buildings*, vol. 43, no. 10, pp. 2918–2925, 2011.

[36] Q. Xu and S. Dubljevic, "Model predictive control of solar thermal system with borehole seasonal storage," *Computers & Chemical Engineering*, vol. 101, pp. 59–72, 2017.

[37] W. Wei, C. Gu, D. Huo, S. Le Blond, and X. Yan, "Optimal borehole energy storage charging strategy in a low carbon space heat system," *IEEE Access*, vol. 6, pp. 76 176–76 186, 2018.

[38] J. Lago, F. De Ridder, P. Vrancx, and B. De Schutter, "Forecasting day-ahead electricity prices in Europe: The importance of considering market integration," *Applied Energy*, vol. 211, pp. 890–903, 2018.

[39] J. Lago, K. De Brabandere, F. De Ridder, and B. De Schutter, "Short-term forecasting of solar irradiance without local telemetry: A generalized model using satellite data," *Solar Energy*, vol. 173, pp. 566–577, 2018.

[40] F. Feijoo, W. Silva, and T. K. Das, "A computationally efficient electricity price forecasting model for real time energy markets," *Energy Conversion and Management*, vol. 113, pp. 27–35, 2016.

[41] Y. Ji, R. J. Thomas, and L. Tong, "Probabilistic forecast of real-time LMP via multiparametric programming," in *2015 48th Hawaii International Conference on System Sciences*, 2015.

[42] J. B. Rawlings, D. Q. Mayne, and M. M. Diehl, *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing, 2017.

[43] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. MIT press, 2018.

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

[45] H. Louie, "Evaluation of bivariate archimedean and elliptical copulas to model wind power dependency structures," *Wind Energy*, vol. 17, no. 2, pp. 225–240, 2014.

[46] F. Golestaneh, H. B. Gooi, and P. Pinson, "Generation and evaluation of space–time trajectories of photovoltaic power," *Applied Energy*, vol. 176, pp. 80–91, 2016.

[47] J. Lago, F. De Ridder, W. Mazairac, and B. De Schutter, "A 1-dimensional continuous and smooth model for thermally stratified storage tanks including mixing and buoyancy," *Applied Energy*, vol. 248, pp. 640–655, 2019.

[48] H. G. Bock and K.-J. Plitt, "A multiple shooting algorithm for direct solution of optimal control problems," *IFAC Proceedings Volumes*, vol. 17, no. 2, pp. 1603–1608, 1984.

[49] D. Ernst, P. Geurts, and L. Wehenkel, "Tree-based batch mode reinforcement learning," *Journal of Machine Learning Research*, vol. 6, pp. 503–556, 2005.

[50] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 785–794.

[51] L. Busoniu, R. Babuška, and B. De Schutter, "A comprehensive survey of multiagent reinforcement learning," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 38, no. 2, pp. 156–172, 2008.

[52] C. Claus and C. Boutilier, "The dynamics of reinforcement learning in cooperative multiagent systems," in *Proceedings of the Fifteenth National/Tenth Conference on Artificial Intelligence/Innovative Applications of Artificial Intelligence*. American Association for Artificial Intelligence, 1998, pp. 746–752.

[53] L. Matignon, G. J. Laurent, and N. Le Fort-Piat, "Independent reinforcement learners in cooperative markov games: a survey regarding coordination problems," *The Knowledge Engineering Review*, vol. 27, no. 1, pp. 1–31, 2012.

[54] Ecovat. (2018) Ecovat renewable energy technologies BV. [Online]. Available: http://www.ecovat.eu

[55] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi – A software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, pp. 1–36, 2018.

[56] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical Programming*, vol. 106, no. 1, pp. 25–57, 2006.

[57] R. Weron, "Electricity price forecasting: A review of the state-of-the-art with a look into the future," *International Journal of Forecasting*, vol. 30, no. 4, pp. 1030–1081, 2014.

[58] K. Høyland and S. W. Wallace, "Generating scenario trees for multistage decision problems," *Management Science*, vol. 47, no. 2, pp. 295–307, 2001.

[59] A. Shapiro and A. Nemirovski, "On complexity of stochastic programming problems," in *Continuous optimization*. Springer, 2005, pp. 111–146.

[60] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*, ser. Springer Series in Statistics. New York, NY, USA: Springer New York Inc., 2001.

[61] J. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl, "Algorithms for hyper-parameter optimization," in *Advances in Neural Information Processing Systems*, 2011, pp. 2546–2554.

[62] Y. Ji, J. Kim, R. J. Thomas, and L. Tong, "Forecasting real-time locational marginal price: A state space approach," in *2013 Asilomar Conference on Signals, Systems and Computers*, 2013.

[63] S. Stoft, *Power System Economics: Designing Markets for Electricity*. Wiley, 2002.

[64] G. Erbach, "Understanding Electricity Markets in the EU," European Parliamentary Research Service, Tech. Rep., 2016.

[65] R. A. van der Veen and R. A. Hakvoort, "The electricity balancing market: Exploring the design challenge," *Utilities Policy*, vol. 43, pp. 186–194, 2016.

[66] G. Pepermans, J. Driesen, D. Haeseldonckx, R. Belmans, and W. D'haeseleer, "Distributed generation: definition, benefits and issues," *Energy Policy*, vol. 33, no. 6, pp. 787–798, 2005.

[67] K. Poplavskaya and L. De Vries, "Distributed energy resources and the organized balancing market: A symbiosis yet? case of three European balancing markets," *Energy Policy*, vol. 126, pp. 264–276, 2019.

[68] R. A. van der Veen, A. Abbasy, and R. A. Hakvoort, "Agent-based analysis of the impact of the imbalance pricing mechanism on market behavior in electricity balancing markets," *Energy Economics*, vol. 34, no. 4, pp. 874–881, 2012.

[69] (2019) Imbalance pricing system: How are the (directions of) payment determined? TenneT TSO B.V. [Online]. Available: https://www.tennet.eu/fileadmin/user_upload/SO_NL/ALG_imbalance_pricing_system.doc.pdf