Optimization and Engineering
https://doi.org/10.1007/s11081-025-10050-9

RESEARCH ARTICLE

Aging-aware battery control via convex optimization

Obidike Nnorom Jr.[1] · **Mehmet Giray Ogut**[1] · **Stephen Boyd**[1] · **Philip Levis**[2]

Received: 14 May 2025 / Revised: 24 August 2025 / Accepted: 18 October 2025

Abstract

We consider the task of controlling a battery while balancing two competing objectives that evolve over different time scales. The primary objective, such as generating revenue by exploiting time varying energy prices or smoothing out the load of a computation center, operates on the scale of hours or days. The long term objective is to maximize the lifetime of the battery, which operates on a time scale of months and years. These objectives conflict; roughly speaking, the primary objective improves with cycling the battery more, which ages the battery faster. Using an existing model for battery aging, we formulate the problem of controlling the battery under these competing objectives as a convex optimization problem. We demonstrate the tradeoff between the primary objective and battery lifetime through numerical simulations.

Contents

1 Introduction
1.1 Setting and tasks
1.2 Related work
1.3 Outline
2 Cell aging model
2.1 Cell charge and current
2.2 Cell aging
2.3 Battery aging model
2.4 Numerical example
3 Arbitrage

✉ Mehmet Giray Ogut
giray98@stanford.edu

Obidike Nnorom Jr.
obdk@stanford.edu

Stephen Boyd
boyd@stanford.edu

Philip Levis
pal@cs.stanford.edu

[1] Department of Electrical Engineering, Stanford University, Palo Alto, USA

[2] Department of Computer Science, Stanford University, Palo Alto, USA

Published online: 19 November 2025

3.1 Problem setup
3.2 Short term MPC method
3.3 Data
3.4 Simulation results
4 Load smoothing
4.1 Problem setup
4.2 Short term MPC method
4.3 Data
4.4 Load forecasts
4.5 Simulation results
5 Conclusion
References

1 Introduction

1.1 Setting and tasks

In the last decade, battery storage, whether retail or grid scale, has become increasingly common. The ability to store energy and discharge it at a later time allows for temporal decoupling of energy production and consumption. Providing this service can be lucrative, as the price of electricity can vary by orders of magnitude depending on the time of day and the season. In other basic applications, a battery can be used to smooth out the power produced by a renewable source, or consumed by a load such as a computation center.

Like any physical system, batteries degrade over time. The more a battery is cycled, meaning charged and discharged, the more it ages, and the shorter its lifetime becomes. The trade-off between short term tasks such as arbitrage and renewable smoothing, and a long term task of maximizing battery lifetime, is the focus of this paper.
Short term objectives. We consider two applications, as simple examples of short term objectives.

1. **Energy arbitrage.** The battery is charged when electricity is cheap and discharged when it is expensive. We evaluate the performance of the system by its revenue.
2. **Load smoothing.** The battery is used to smooth out the electric power demand (*i.e.*, load) consumed by a computation center, so the net power varies more smoothly over time. We evaluate the performance of the system using the root mean square difference between the load and its previous value.

These are just two simple illustrative examples of short term objectives; many others could be handled by the methods we describe in this paper.

Long term objective. The battery has a finite lifetime, which is determined primarily by the number of cycles it undergoes. We want to maximize the lifetime of the battery, defined as the time the battery capacity drops below some fraction of its initial value such as 80% or 90%. Without considering battery aging, the two short term tasks described above involve aggressive cycling of the battery, leading to a shortened lifetime. We show how to operate the battery so as to achieve an optimal trade-off of the short term objective and the long term objective.

1.2 Related work

Battery aging models. Battery aging is often split into two main effects: calendar aging, which happens when the battery is resting (no charge or discharge), and cycle aging, which happens during active charge–discharge cycles (Vermeer et al. 2021). Both processes are sensitive to temperature (especially above 30°C), high current rates, and usage patterns like depth of discharge or state of charge (SoC) (Xiong et al. 2020).

When it comes to modeling these effects, three main approaches appear in the literature. Electrochemical models (Keil and Jossen 2020; Li et al. 2020; Allam and Onori 2020) try to describe the internal reactions mathematically (for instance, using the Butler–Volmer equations (Latz and Zausch 2013)) and can be quite detailed but hard to implement in practice. At the other extreme, empirical models fit observed data to capture aging trends, but these can fail outside their specific test conditions and often need large datasets (Pelletier et al. 2017). Semi-empirical models (Serrao et al. 2009; Suri and Onori 2016; Marano et al. 2009; Torregrosa et al. 2024) blend theoretical ideas with curve-fitting so that the most critical aging drivers (like temperature, SoC, or C-rate) are handled without the complexity of a full electrochemical model or the narrow scope of an entirely data-driven approach.

Calendar aging mainly depends on how batteries are stored (in terms of temperature and SoC), with slow chemical reactions gradually eating away at capacity. Cycle aging, on the other hand, is tied to how often and how aggressively the battery is charged or discharged. In actual operation, both processes happen at the same time, so an aging model typically combines or overlays both (Liu et al. 2020). Semi-empirical models are especially attractive in this setting because they balance realism and simplicity, letting us capture key aging behaviors in a way that can be used in practical control algorithms (Jin et al. 2018; Miller et al. 2022).

Optimization in batteries. Current optimization techniques used in battery longevity do not directly focus on working with aging models of the batteries. Liu et al. (2018) focuses on optimizing charging behavior of both CC-CV charging and multi stage CC-CV charging. Nonlinear optimization techniques are used to determine the optimal current and voltage values to balance between aging, efficiency loss, and charge time. Chung et al. (2020) focuses on reducing the calendar aging of a PEV battery through an optimal charging scheme. This paper argues that calendar aging is most important for PEV batteries, and proposes a nonlinear interior point method to determine an optimal overnight charging scheme for the PEV battery.

Bashir et al. (2017) tackles lifetime maximization of lead-acid batteries by formulating key aging characteristics of the battery as convex formulas. Bad recharge, the time since last full charge, and the lowest state of charge since last recharge are the three primary factors in their objective function. A multi-objective convex optimization problem is then solved to maximize the lifetime of these batteries.

Aging aware optimization methods. The adoption of batteries as a power source in many systems has led to increased focus on degradation. Researchers have focused on characterizing the aging of battery cells through a variety of different aging models. Simple approaches model battery aging based on a few factors such as depth of discharge (DOD) while more advanced models, such as physics based models, cap-

ture internal degradation mechanisms through coupled differential equations (Franco et al. 2016). Simple models are inadequate for accurately representing battery aging in most of the applications in which aging is relevant to consider while the complexity of physics based models make them hard to use in real-time applications (Franco et al. 2016). Semi-empirical models, on the other hand, are guided by physics but are simplified. As a result, these models can still accurately capture aging behavior while working more seamlessly under real-time applications.

The consideration of battery aging has increasingly become a concern during long term battery control. Optimization techniques such as Mixed Integer Linear Programs (MILPs) are frequently used for aging-aware BESS scheduling in the literature. Maheshwari et al. developed a nonlinear Li-ion degradation model from experiment data and cast it as an MILP (Maheshwari et al. 2020). Xu et al. constructs a convex function to represent aging but uses binary variables to determine when the battery charges and discharges causing the entire optimization problem to be an MILP (Xu et al. 2018). Although MILPs are popular in the literature, solving an MILP is an NP hard problem and thus the solve time can scale exponentially with the number of variables in the problem (Garey and Johnson 1990; Collath et al. 2022).

Dynamic programming (DP) and related decomposition methods are classical tools for sequential decision problems and have been used in battery scheduling. Abdulla et. al employs a stochastic DP for a PV-battery system, where the immediate cost function includes a penalty for capacity loss and the state transition accounts for degradation (Abdulla et al. 2018). The paper reports a significant extension (up to 160%) of battery lifetime through this technique compared to naive approaches. Dynamic programs are capable of finding true optimal policies for a given model, but are not inadequate for problems with long time horizons (Bellman 1957). Exact DP becomes computationally expensive as the number of timesteps or state granularity increase. This paper gets around this limitation by assuming there are a finite number of discrete SOCs possible which is not the case in practice.

Besides the two techniques mentioned above, numerous hueristics based models have been applied to aging-aware dispatch. These heuristics include rule-based strategies (e.g. simple rules like "avoid discharge below 20% SOC). Bashir et. al uses this method effectively in their work by identifying the three most important factors that lead to aging based on the Shiffer model for lead acid batteries (Bashir et al. 2017). The authors are able to take these rules and represent as convex functions that they plug into their model. Weitzel et al. provides an analysis of heuristic based models for aging (Weitzel and Glock 2018). Heuristic models are great for capturing large trends through simple functions, but don't fully capture aging like a physics based model or semi-empirical model can.

Model predictive control (MPC). Model predictive control goes by several other names, such as rolling-horizon planning, receding-horizon control, dynamic matrix control, and dynamic linear programming. Originally developed in the 1960s, MPC offers a powerful framework for managing constraints on states, inputs, and outputs. It has a long history and large literature, and is widely used. Some early work is Cutler (1979); Garcia et al. (1989); for more recent surveys see the papers (Holkar and Waghmare 2010; Mayne 2014; Abughalieh and Alawneh 2019; Schwenzer et al. 2021)

or books (Camacho et al. 2007; Borrelli et al. 2017; Grüne et al. 2017; Rawlings et al. 2017; Rakovic and Levine 2018).

Papers describing applications of MPC in specific areas include HEVs (Huang et al. 2017) data center cooling (Lazic et al. 2018), building HVAC control (Afram and Janabi-Sharifi 2014), wind power systems (Hovgaard et al. 2015), microgrids (Hu et al. 2021), pandemic management (Carli et al. 2020; Péni et al. 2020), dynamic hedging (Primbs 2009), revenue management (Talluri and Ryzin 2006; Bertsimas and Popescu 2003), railway systems (Felez et al. 2019), aerospace systems (Eren et al. 2017), and agriculture (Ding et al. 2018). With appropriate forecasting (which in many applications is typically simple) and choice of cost function, MPC can work well, even though it does not explicitly take into account uncertainty in the dynamics and cost, or more precisely, since it is based on a single forecast of these quantities.

There are many extensions of MPC that attempt to improve performance by taking into account uncertainty in the future dynamics and cost. Examples include robust MPC (Bemporad and Morari 2007; Campo and Morari 1987), min-max MPC (Raimondo et al. 2009), tube MPC (Mayne et al. 2005), stochastic MPC (Heirung et al. 2018; Mesbah 2016) and multi-forecast MPC (MF-MPC) (Shen and Boyd 2021).

A widely recognized shortcoming of MPC is that it can usually only be used in applications with slow dynamics, where the sample time is measured in seconds or minutes. However there exist methods to speed up MPC, such as computing the entire control law offline or using online optimization (Wang and Boyd 2009).

Since it integrates constraint handling, future forecasting, and feedback adjustment, MPC is often viewed as a middle ground between exhaustive search methods like dynamic programming (DP) and simpler real-time strategies such as the equivalent consumption minimization strategy (ECMS) (Zhang et al. 2015). This balance between computational tractability and robust performance has made MPC increasingly popular in both academic and industrial settings for battery management systems.

Energy arbitrage. In economics and finance, arbitrage is the practice of taking advantage of a price difference by buying energy from the grid at a low price and selling it back to the grid at a higher price (Zafirakis et al. 2016). Although it is often assumed to occur in the day-ahead markets (Staffell and Rustomji 2016; Wilson et al. 2018), arbitrage strategies for intraday markets have also been considered (Metz and Saraiva 2018). Multiple studies assessed how to maximize arbitrage profits (Sioshansi et al. 2009), but the consistent finding is that the attainable revenues are on their own insufficient to repay investment in battery storage. Bradbury et al. (2014) showed that Li-ion batteries often fail to surpass 0% IRR in U.S. markets, though short charge times, lower capital costs, and ancillary services could enhance returns. Sioshansi et al. (2009) found that arbitrage value in PJM depends on round-trip efficiency, location, and fuel mix, noting that lower natural gas prices warrant re-evaluation. McConnell et al. (2015) highlighted possible profitability in 5-minute dispatch markets and potential peaker displacement, suggesting that as renewable penetration depresses prices, arbitrage opportunities may grow in the United States.

The economic value that a battery operator can obtain from arbitrage rests on both technical and market elements. Among the technical considerations, round-trip efficiency stands out for its influence on marginal operating costs (Critchlow and

Denman 2017). Another crucial factor is the discharge capacity, or energy-to-power ratio, which determines how much energy a battery can store. Because physical and operational stresses cause capacity fade over the asset's lifetime (Schmidt et al. 2017), they significantly affect profitability (He et al. 2020). Although the replacement cost of a battery is typically incurred only at the end of its lifespan, Xu et al. (2017) notes that aging–driven by operational choices–should still factor into marginal cost calculations, as it can alter the operating strategy itself. On the market side, price volatility rather than average price levels is generally acknowledged as the main determinant of arbitrage value (Wilson et al. 2018).

Renewable generation and load smoothing. A key challenge for large-scale renewable integration is the inherent fluctuation in power generation, which can cause frequency deviations, voltage inconsistencies, and high peak loads. A common solution is to pair wind turbines with batteries to smooth out these variations, store surplus energy during periods of high generation, and feed it back when generation drops (Díaz-González et al. 2013).

The issues caused by wind power fluctuations were first discussed in the literature in the early 1980s, when commercial wind turbines started being installed more regularly. In the first studies, the authors proposed less sophisticated methods of power smoothing (Suvire et al. 2012). In the late 1990s, however, more studies began to consider storage systems (mainly fly-wheel and lead-acid batteries) to smooth the output power from wind turbines (Jerbi et al. 2009; El-Naga et al. 2017; Elkomy et al. 2017).

In 2009, Khalid and Savkin (2009, 2010) presented a controller design for wind power smoothing purposes based on model predictive control. They noted that prediction could help improve the economy and security of wind integration into electrical grids. Thus, a wind power prediction system combined with a battery was proposed based on measurements from different observation points and communication channels. The effectiveness of this approach was assessed through real wind speed data from an Australian wind farm comprising 37 wind turbines. The results show the capability of the controller to smooth the wind power, optimize the maximum ramp rate requirement, and also the state of charge of the battery. The study accounted for inefficiencies in batteries in terms of energy conversion but did not consider the battery aging.

An analogous problem tackles the challenge of smoothing a load that varies rapidly over time, such as one that might appear in a data computation center while processing a job such as training a large language model (Li et al. 2024). The existing literature explores the use of energy storage that is either integrated in datacenter uninterruptible power supply (UPS) systems or deployed as standalone battery banks to enable demand response services (Urgaonkar et al. 2011; Govindan et al. 2011, 2012; Mamun et al. 2015). Most studies concentrate on minimizing the total cost of ownership (TCO) (Wang et al. 2014; Kontorinis et al. 2012), defined as the sum of amortized capital expenditures and operating costs over a prescribed time horizon (Barroso et al. 2019). Queueing-theoretic Lyapunov optimization has also been used to derive policies that nearly minimize monthly electricity bills (Urgaonkar et al. 2011; Guo et al. 2011). While these analyses predominantly assume lead-acid batteries, a subset of work considers lithium-ion technology. In such demand response

formulations, lithium-ion units are typically modelled as ideal charge integrators, and aging is captured through depth-of-discharge (DoD) charts or charge throughput heuristics (Kontorinis et al. 2012; Wang et al. 2012; Aksanli et al. 2013; Ren et al. 2012).

In 2016, Mamun et al. (2016) formulated a multi-objective framework for datacentre demand response that couples a nonlinear equivalent circuit model with SEI based aging and tunes feedforward feedback controllers on a lithium-ion battery pack. Their results reveal an inherent tradeoff between cost savings and battery health; dead-band PI control mitigates this compromise and remains robust to load uncertainty and battery pack size.

1.3 Outline

We describe the battery model in §2 where we use an existing semi-empirical aging model from Serrao et al. (2009); Suri and Onori (2016) and come up with a convex approximation of the aging rate. Next, we describe two applications in §3 and §4 where we use model predictive control for price arbitrage and load smoothing and give numerical examples.

2 Cell aging model

This section explains battery aging, a mathematical model of how a single battery cell ages (Serrao et al. 2009; Suri and Onori 2016), and how we use this model to predict the aging of a larger, multi-cell battery. To distinguish these two cases we refer to the former as a "cell" and the latter as a "battery".

We consider a lithium iron phosphate ($LiFePO_4$)-graphite battery. Cells operate by storing chemical potential energy. When a voltage is applied across the battery terminals, the electrical field causes ions to move through the battery's electrolyte, converting the electrical potential into chemical potential. This process decays the material of the cell, especially the cathode and anode. They can crack and oxidize; films can form on them; ions can become embedded in them. There are complex and detailed physical models for this process, as predicting lifetime is important in battery management.

2.1 Cell charge and current

We use Suri et al.'s model of lithium-iron phosphate cell aging (Suri and Onori 2016). This model is for a single, 2.5 Ampere-hour (Ah), 3.3 Volt (V) lithium-ion cell. Physical models such as these model a cell using charging current in (A) and model capacity and charge in (Ah). When we model a multi-cell battery, we switch to the more convenient power and energy units, Watts (W) and Watt-hours (Wh).

We model the cell charging current as constant over time intervals of length δ hours, so, *e.g.*, $\delta = 0.25$ means 15 minute intervals. We denote the time periods as $t = 1, 2, \ldots$. We denote the (instantaneous) charge in the cell at the beginning of

interval t as $\tilde{q}_t$, in units of (Ah). The charge satisfies $0 \leq \tilde{q}_t \leq \tilde{Q}_t$, where $\tilde{Q}_t > 0$ is the capacity in (Ah) of the cell at time in interval t. We assume that $\tilde{Q}_t$ is known (or measured) at time t. We refer to $\tilde{Q}_1$ as the initial cell capacity, and $\tilde{Q}_t$ as the cell capacity at time t. The values $\tilde{q}_t = 0$ and $\tilde{q}_t = \tilde{Q}_t$ mean that the cell is empty and full, respectively. The empty cell charge $\tilde{q}_t = 0$ refers to the lowest charge of the cell over its useful range, and not absolute zero cell charge. Similarly, $\tilde{q}_t = \tilde{Q}_t$ refers to the largest charge of the cell over its useful range.

The cell (discharge) current in interval t is denoted $\tilde{b}_t$, in (A). Positive values of $\tilde{b}_t$ correspond to discharging the cell, and negative values of $\tilde{b}_t$ correspond to charging the cell. The cell current $\tilde{b}_t$ must satisfy $|\tilde{b}_t| \leq \tilde{B}$, where $\tilde{B} > 0$ is the maximum cell charge and discharge current in (A), given by $\tilde{B} = \tilde{Q}_1 C$, where $C > 0$ is the maximum C-rate of the cell, in inverse hours (1/h). The cell dynamics are given by $\tilde{q}_{t+1} = \tilde{q}_t - \delta \tilde{b}_t$. Note that $\delta \tilde{b}_t$ is the total charge, in (Ah), removed from the cell in interval t.

Short and long term quantities. We refer to $\tilde{b}_t$ and $\tilde{q}_t$ as short term quantities, since they can vary considerably from interval to interval. We refer to $\tilde{Q}_t$ as a long term (aging) quantity, since it changes very slowly, with appreciable change only over a time period measured in months or longer. In particular, $\tilde{Q}_t$ can be considered approximately constant over a period on the order of days.

2.2 Cell aging

Loss and loss rate. As the cell is used its capacity $\tilde{Q}_t$ decreases, *i.e.*, $\tilde{Q}_{t+1} \leq \tilde{Q}_t$. The lifetime L of the cell in (h) is the time L when the cell capacity drops below some fixed fraction of its initial value, such as 90%, *i.e.*, $\tilde{Q}_{L-1} \geq 0.9 Q_1$ and $\tilde{Q}_L < 0.9 Q_1$. (It is also common to define lifetime using 80% of the initial capacity.) We define the normalized capacity loss as $l_t = \left(Q_1 - \tilde{Q}_t\right)/Q_1$, which we can express as a percentage, so, *e.g.*, $l_t = 0.07$ means the cell has experienced 7% capacity loss. The normalized capacity loss starts at $l_1 = 0$ (no capacity loss), and rises until $l_t > 0.1$, corresponding to cell lifetime.

The cell aging rate is defined as $\rho_t = (l_{t+1} - l_t)/\delta$, so we have

$$l_t = \delta \sum_{\tau=1}^{t-1} \rho_\tau.$$

The aging rate gives the increase in normalized capacity loss per hour, and has units (1/h). A cell with a constant aging rate ρ has a lifetime around $L = 0.2/\delta$ hours, which is $2.28 \times 10^{-5}/\delta$ in years, a more commonly used time unit for cell lifetime. Typical values of ρ_t are on the order of 10^{-6} or 10^{-5}, corresponding to cell lifetime ranging from around 2 to 20 years.

Loss rate model. The aging rate ρ_t depends on how the cell is used, *i.e.*, its history of charging and discharging up to period t. In this paper we use the semi-empirical aging model for a lithium iron phosphate ($LiFePO_4$)-graphite cell frequently used in hybrid electric vehicle batteries, given in Suri and Onori (2016). (But our methods can be used with any other specific cell aging model.) The model is

$$\rho_t = z\left(\sum_{\tau=1}^{t} |\tilde{b}_\tau|\delta\right)^{z-1} |\tilde{b}_t| \left(\alpha \frac{\tilde{q}_t}{\tilde{Q}_t} + \beta\right) \exp\left(\frac{-E_a + \eta \frac{|\tilde{b}_t|}{\tilde{Q}_t}}{R_g T}\right). \tag{1}$$

Here $\sum_{\tau=1}^{t} |\tilde{b}_\tau|\delta$ is the accumulated (absolute) charge throughput, up to time t, with units (Ah). We note that $\tilde{q}_t/\tilde{Q}_t$ is the cell charge normalized to its capacity (between 0 and 1), and $|\tilde{b}_t|/\tilde{Q}_t$ is the instantaneous C-rate of the cell, which is between 0 and C.

The terms and constants appearing in (1) are as follows.

- The first term models the effect of accumulated charge on aging. The (unitless) power law exponent is $z = 0.60$.
- The second term models the effect of instantanous cell charge of aging rate. We take $\alpha = 28.966$ and $\beta = 74.112$, with units $(1/\text{Ah}^z)$.
- The last term comes from the Arrhenius equation and gives the dependence of aging on temperature and instantaneous charging rate. Here $E_a = 31500$ is the activation energy with units (Jmol^{-1}), $R_g = 8.314$ is the universal gas constant with units $(\text{Jmol}^{-1}\text{K}^{-1})$, and $\eta = 152.500$, with units $(\text{Jmol}^{-1}\text{h})$. The absolute temperature T is given in degrees Kelvin (K).

The accumulated charge and cell capacity $\tilde{Q}_t$ are long term quantities, which do not change much over periods of a few days (after the initial few months, in the case of the accumulated charge). Combining the long term quantities we can rewrite the capacity loss rate (1) as a function of the short term quantities as

$$\rho_t = \mu_t |\tilde{b}_t|(1 + \nu_t \tilde{q}_t) \exp \lambda_t |\tilde{b}_t|, \tag{2}$$

where

$$\mu_t = \beta \exp\left(\frac{-E_a}{R_g T}\right) z \left(\sum_{\tau=1}^{t} |\tilde{b}_\tau|\delta\right)^{z-1}, \qquad \nu_t = \frac{\alpha}{\beta \tilde{Q}_t}, \qquad \lambda_t = \frac{\eta}{R_g T \tilde{Q}_t}$$

are long term, slowly varying quantities, known at time t. Over short time periods (*e.g.*, a few days), the coefficients μ_t, ν_t, and λ_t can be considered constant, so the aging loss rate depends only on $|\tilde{b}_t|$ and $\tilde{q}_t$. It increases with $|\tilde{b}_t|$ proportional to $|\tilde{b}_t| \exp \lambda_t |\tilde{b}_t|$, and proportional to $1 + \nu_t \tilde{q}_t$ with $\tilde{q}_t$.

Short term approximation. We develop here a convex approximation of ρ_t in (2) that can be used in the short term. Taking the first order Taylor expansion of (2) with respect to $|\tilde{b}_t|$ and $\tilde{q}_t$, around the point $|\tilde{b}_t| = 0$ and $\tilde{q}_t = \tilde{Q}_t/2$, we obtain

$$\hat{\rho}_t = \mu_t \left(1 + \nu_t \frac{\tilde{Q}_t}{2}\right) |\tilde{b}_t|. \tag{3}$$

This approximation seems crude, but given the recommended C-rate usage of the battery and our current applications, the Taylor expansion chosen was more than sufficient to accurately represent the semi-empirical aging model. A Taylor expansion could be

taken about other points, but the points chosen allow for a convex approximation of the semi-empirical model that accurately represent the true aging model.

The convex approximation technique used in this work is tailored for models that use the Arrhenius equation to represent aging as shown in (1) and extensively in the literature. The methods to convexify the formula, however, can be applied to a variety of semi-empirical aging models to represent these models in a form that works within a convex framework.

2.3 Battery aging model

The above model is for a single battery cell. Larger batteries are made up of many cells. We consider a multi-cell battery consisting of N cells, either passively wired in a series-parallel arrangement, or with active battery power control, or any combination. We model the multi-cell battery with capacity given in energy (Wh) and charge/discharge given in power (W). We assume that the battery cells work in a balanced way, *i.e.*, all have the same current, charge, and capacity at each time.

We denote the cell charging current as $\tilde{b}_t$ in (A), cell charge as $\tilde{q}_t$ in (Ah), and cell capacity as $\tilde{Q}_t$ in (Ah), as above. We denote the battery charging power in (W) as b_t, the battery charge as q_t in (Wh), and the battery capacity as Q_t in (Wh). These are related as

$$b_t = 3.3N\tilde{b}_t, \quad q_t = 3.3N\tilde{q}_t, \quad Q_t = 3.3N\tilde{Q}_t,$$

where 3.3 (V) is the cell voltage. (Here we make the reasonable approximation that the cell voltage is constant over its useful charge range.) The maximum C-rate of the battery is the same as that of a single cell.

2.4 Numerical example

In this section we demonstrate the aging model (2) and the short-term approximation (3) with examples. The battery has $N = 5000000$ cells, corresponding to a capacity $Q_1 = 4.125$ (MWh). The maximum C-rate is set to 0.33 (1/h) which means the battery can completely charge or discharge in 3 hours. We consider two simple charging profiles, shown in Figure 1. The first one fully charges and discharges the battery, with constant charge/discharge current, with 2 cycles per day. The second one does the same, with 4 cycles per day.

The associated capacity losses are shown in Figure 2. In these plots we show the aging capacity loss calculated using the exact aging model (2) and the approximate aging model (3). The approximate aging model matches the exact aging model well: Table 1 shows it predicts lifetimes within a few percent of the more accurate model.

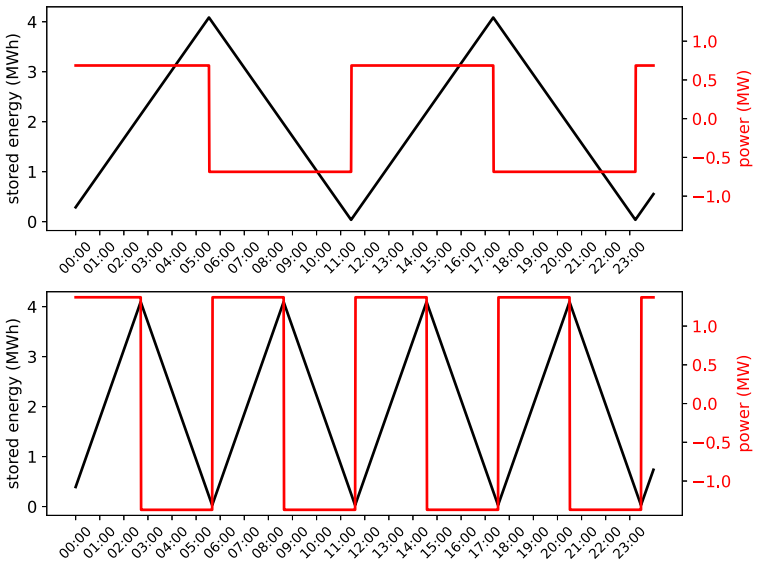

Fig. 1 Charging profiles. *Top.* 2 cycles per day. *Bottom.* 4 cycles per day

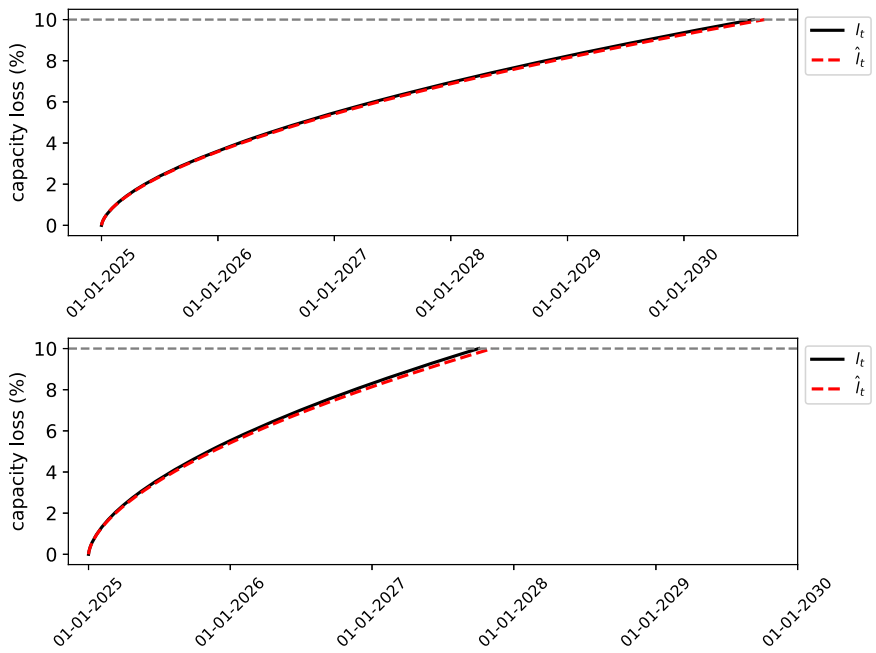

Fig. 2 Actual and approximate aging. *Top.* 2 cycles per day. *Bottom.* 4 cycles per day

Table 1 Lifetime for different charging profiles, calculated using exact and the approximate aging models

Charging profile (cycles/day)	Actual lifetime (years)	Approximate lifetime (years)
2	5.60	5.70
4	2.75	2.85

3 Arbitrage

3.1 Problem setup

In our first example the short term task is arbitrage with time-varying energy prices. The (nonnegative) price of electricity in period t is p_t in (USD/MWh). The battery is connected to the grid, and a payment of $p_t b_t \delta$ is received in period t; negative payments are amounts we pay to the grid operator. The objective is to choose b_t to maximize the average (net) payment over a day given by $\frac{1}{T^{\text{day}}} \sum_{\tau=1}^{T^{\text{day}}} p_\tau b_\tau \delta$, where T^{day} is the number of time periods in a day and δ is the period length in hours (h). The prices are known far enough ahead of time (*e.g.*, on day), that we can consider them known.

Without considering battery aging, the greedy strategy is to discharge the battery as much as possible when the price is high, and charge it as much as possible when the price is low. In that case, we are only subject to physical limits such as maximum discharge rate, battery capacity, current storage, etc. However this strategy will lead to aggressive cycling of the battery and will shorten its lifetime. That is why we also consider the total revenue over the lifetime of the battery, $\sum_{\tau=1}^{T^{\text{EOL}}} p_\tau b_\tau \delta$, where T^{EOL} is the end of life of the battery.

3.2 Short term MPC method

At time t we are given battery storage q_t, battery capacity Q_t, and compute the approximate aging rate coefficient $\mu_t(1 + \nu_t \frac{Q_t}{2})$ using the battery model in §2. We consider a horizon of H periods, and assume that together with Q_t, the aging rate coefficient is constant over this horizon. We assume that the price of electricity p_t is known for the next H periods and that they are nonnegative. To find the optimal battery discharge b_τ, and storage q_τ for $\tau = t+1, \ldots, t+H$, we solve the problem

$$\begin{array}{ll} \text{minimize} & \frac{1}{H} \sum_{\tau=t+1}^{t+H} \left(-p_\tau b_\tau \delta + \gamma \mu_t (1 + \nu_t \frac{Q_t}{2}) |b_\tau| \right) + \eta \left(q_{t+H} - \frac{Q_t}{2} \right)^2 \\ \text{subject to} & q_\tau = q_{\tau-1} - \delta b_{\tau-1}, \quad \tau = t+1, \ldots, t+H \\ & |b_\tau| \le C Q_t, \quad \tau = t+1, \ldots, t+H \\ & 0 \le q_\tau \le Q_t, \quad \tau = t+1, \ldots, t+H \end{array}$$

where $\gamma > 0$ is a trade-off parameter between revenue and battery aging, and $\eta > 0$ is a parameter that penalizes deviation of the terminal battery charge from $Q_t/2$, half

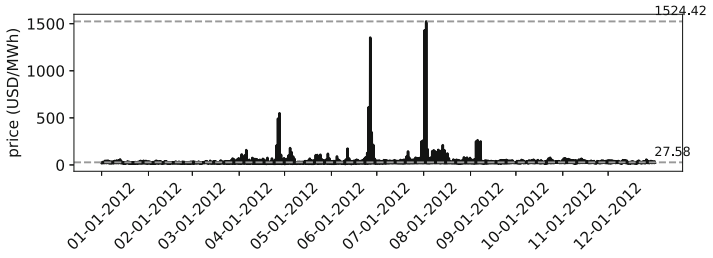

Fig. 3 ERCOT hourly local marginal prices for day-ahead market in 2012. The dashed lines show the mean and the maximum

capacity. Solving this optimization problem gives us a plan for operating the battery over the next H periods. Our policy uses the first battery power b_t in our plan.

We use the open source package CVXPY (Diamond and Boyd 2016) to formulate the problem, and the specific solver CLARABEL (Goulart and Chen 2024) to solve the problems.

3.3 Data

We use hourly local marginal prices (LMP) data for the day-ahead market in the Electric Reliability Council of Texas (ERCOT) North Hub for the year 2012. The data is available at ERCOT (2012). We use data from January 1, 2012 to December 31, 2012 and have 8784 data points. We run simulations until the battery end of life, which ranges from 6 to 22 years.

To simulate multiple years, we repeat the data from 2012, removing the data from February 29th for non-leap years. Repeating the same data for multiple years removes yearly variations in the data, which varies considerably from year to year. For instance, in February 2021, the price of electricity in Texas spiked to $\sim$ 9000 (USD/MWh), two orders of magnitude higher than the average price in 2012, due to a winter storm. To reduce the effect of such outliers and yearly variations, we use the same data for multiple years.

To visualize seasonal variation, we first look at prices at the yearly scale, shown in Figure 3. The mean is 27.58 (USD/MWh) and its standard deviation is 39.36 (USD/MWh). The maximum price is 1524.42 (USD/MWh). We observe that there are 2 major peaks, both in summer months.

Next, we look at the data at a finer scale. Figure 4 shows the prices for the first week of January 2012 and the first week of June 2012. We observe several expected phenomena, *e.g.*, prices are higher during the day than at night, and a bit higher in winter than in summer. We can also see some small variation over a week. One interesting observation is that the shape of the daily demand in winter differs considerably from the shape of the daily demand in summer. In winter we see a double bump, with peaks in the morning and afternoon, while in summer we see a smoother daily variation with one peak in the early afternoon. We interpret this as the effect of heating in winter, which causes a peak in the morning, and air conditioning in summer, which causes a peak in the afternoon. Also, even if the daily demand curves are similar for the same

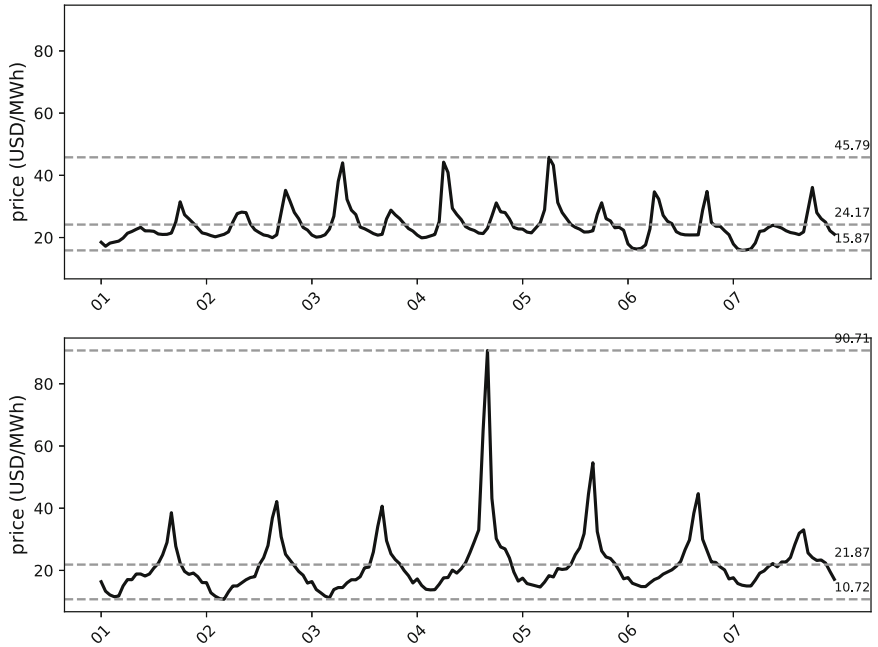

Fig. 4 Two different weeks of price data. *Top.* January 2012. *Bottom.* June 2012

season, they show some variation. For instance, the daily demand on 06-04-2012 is higher with a sharper peak around noon compared to the other days of the week.

We assume that we know the prices for the next H periods beforehand regardless of the time of day, so we do not need to rely on a forecasting model to predict the prices. However, in practice, day-ahead market prices in ERCOT are only available at around 1:00 PM for 00:00 AM to 11:00 PM on the next day. As a result, we would need to use a forecaster for at least some of the hours.

3.4 Simulation results

Using the MPC method in §3.2 and data in §3.3, we change the cost of battery aging γ and simulate hourly (*i.e.*, $\delta = 1$ (h)) until the end of life of the battery. We take horizon $H = 24$, *i.e.*, one day. We found that the results are not sensitive to the value of η, and use $\eta = 1$. Each subproblem takes on the order of milliseconds to solve.

The total revenue and average revenue versus battery lifetime is shown in Figure 5. We observe that the total revenue increases with increasing battery lifetime while the average revenue decreases. This is expected, since with a shorter battery lifetime, the battery is allowed to cycle more aggressively, which increases the hourly revenue. However, this comes at the cost of a shorter battery lifetime and lower total revenue. Hence there is a trade-off between the total revenue and the average revenue.

Since total revenue is a monotonically increasing function of battery lifetime, and we would get higher revenue with a longer battery lifetime, we focus on the net

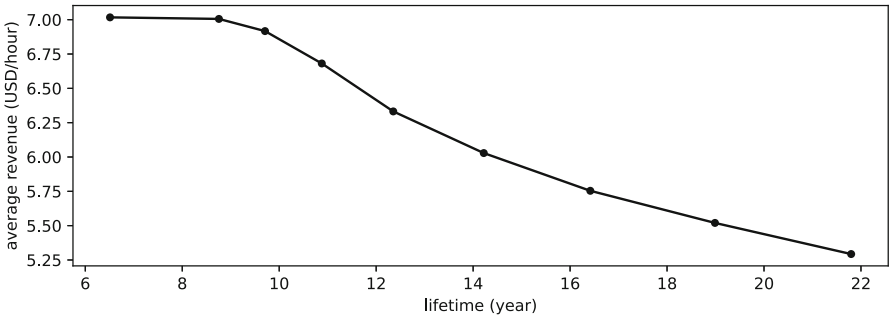

Fig. 5 Average hourly revenue versus battery lifetime

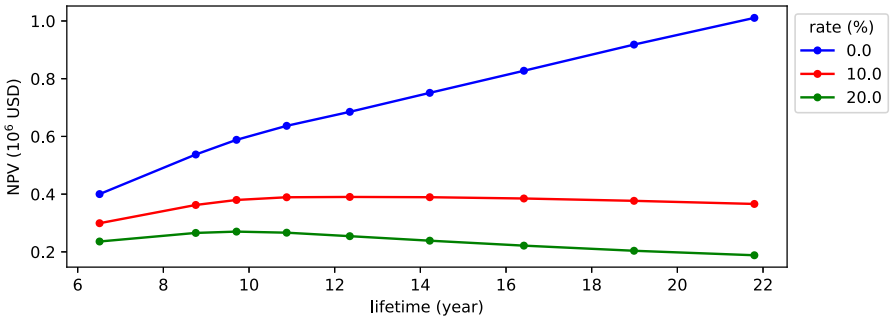

Fig. 6 Net present value of revenue versus battery lifetime

present value (NPV) of the revenue over the lifetime of the battery. The NPV is a discounted sum of the revenue over the lifetime of the battery, where the interest rate corresponds to the cost of capital. For a given interest rate i, NPV is given by $\mathrm{NPV}(i) = \sum_{\tau=1}^{T^{\mathrm{EOL}}} \frac{p_\tau b_\tau \delta}{(1+i)^\tau}$.

We plot the NPV versus battery lifetime in Figure 6 and show for interest rates $i = 0\%, 10\%, 20\%$. We observe that with 20% interest rate, the NPV is maximized with a battery lifetime of around 10 years.

Next, we focus on discharge profiles at a smaller timeframe and look at the weekly discharge profile for a battery with an 8.75 year lifetime in Figure 7. As expected, we observe that discharge periods coincide with high price periods, and charge periods coincide with low price periods. In fact, during the summer the battery is discharged almost every day around noon, where the price is highest, and charged in the early morning, where the price is lowest. In winter, the battery is idle most of the time, and is discharged/charged only a few times a week.

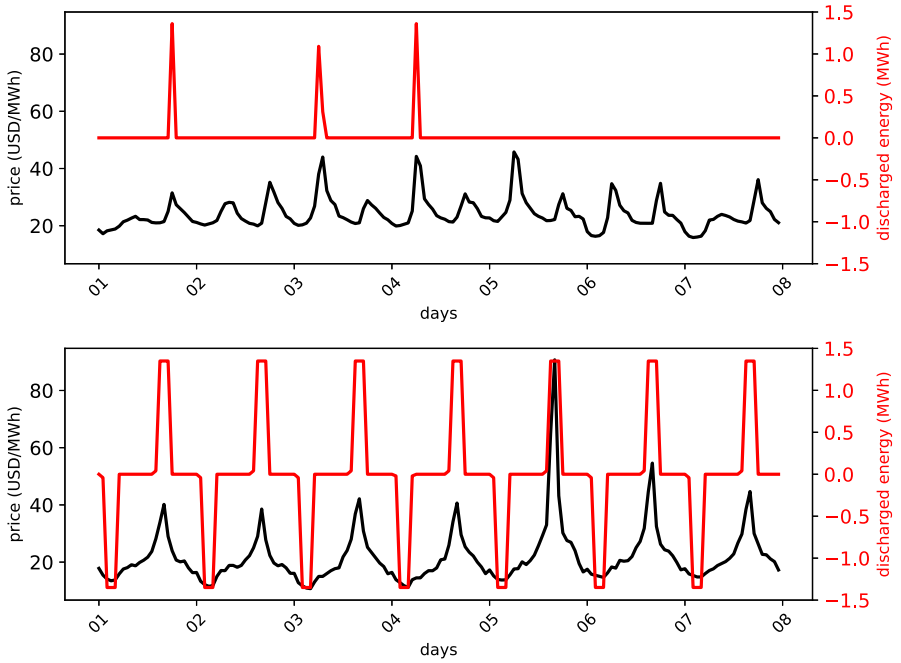

Fig. 7 Weekly discharge profile for a battery with 8.75 years lifetime. *Top.* January 2012. *Bottom.* June 2012

4 Load smoothing

4.1 Problem setup

In our second example the short term task is to smooth out a time-varying load. We have a load w_t in Watts (W), and we operate the battery in parallel, so the total power is $z_t = w_t + b_t$. We want z_t to be smooth, as judged by the RMS difference

$$\mathcal{D} = \left(\frac{1}{T-1} \sum_{t=1}^{T-1} (z_{t+1} - z_t)^2 \right)^{1/2},$$

with smaller values better. At time period t, the load w_t is known; future values $z_{t+1}, z_{t+2}, \ldots$ are not known, but can be forecasted.

4.2 Short term MPC method

We use an MPC policy. At time t, we are given battery storage q_t, battery capacity Q_t, load w_t, and aging parameter $\mu_t \left(1 + \nu_t \frac{Q_t}{2}\right)$. We assume the battery capacity and aging parameter are constant over our short term horizon H. We also have load forecasts $\hat{w}_{\tau|t}$ for $\tau = t+1, \ldots, t+H$. We will discuss how we obtain these forecasts

later in §4.4. We solve the short term planning problem

$$\begin{array}{ll}\text{minimize} & \frac{1}{H}\sum_{\tau=t+1}^{t+H}\left((z_\tau - z_{\tau-1})^2 + \gamma\mu_t(1+\nu_t\frac{Q_t}{2})|b_\tau|\right) + \eta\left(q_{t+H} - \frac{Q_t}{2}\right)^2 \\ \text{subject to} & z_\tau = \hat{w}_{\tau|t} + b_\tau, \quad \tau = t+1,\ldots,t+H \\ & z_t = w_t + b_t \\ & 0 \le z_\tau, \quad \tau = t+1,\ldots,t+H \\ & q_\tau = q_{\tau-1} - \delta b_{\tau-1}, \quad \tau = t+1,\ldots,t+H \\ & |b_\tau| \le CQ_t, \quad \tau = t+1,\ldots,t+H \\ & 0 \le q_\tau \le Q_t, \quad \tau = t+1,\ldots,t+H\end{array}$$

where $\gamma > 0$ is a trade-off parameter between smoothing and battery aging, and $\eta > 0$ penalizes deviation of the final battery storage from half capacity. The solution of this problem gives us a plan for $b_t, \ldots, b_{t+H}$; we use as b_t the first battery charge in this plan.

4.3 Data

We use a battery with $N = 15000$ cells, corresponding to 123.75 (kWh) initial battery capacity. Given that most commercial residential batteries have a capacity of 13.5 (kWh), this is equivalent to having 9 batteries. Maximum charge/discharge rate is $C = 0.3$ (1/h) which implies that the battery can completely charge/discharge in around 3 hours. We use simulated data modeled after the consumption of a large language model (LLM) training job in the MIT Supercloud Dataset as described in Li *et al.* (2024, Figure 6). We use 20 minute periods, and three discrete values of load: 5 (kW), 20 (kW), and 35 (kW), corresponding to different computation states. We refer to these as the low, medium, and high power states, respectively. We use a Markov model for the computation states, taken (roughly) from (2024, Figure 6), with transition matrix

$$P = \begin{bmatrix} 0.79 & 0.22 & 0.00 \\ 0.05 & 0.72 & 0.40 \\ 0.16 & 0.06 & 0.60 \end{bmatrix}$$

where P_{ij} is the probability of transitioning from state j to state i. We generate 25 years of data starting from 01-01-2018 00:00 PST, which gives 657000 data points. The asymptotic state probabilities are 0.40, 0.38, and 0.22. The asymptotic average load power is 17.25 (kW). Figure 8 shows the simulated load data for January 1st 2020.

4.4 Load forecasts

We use a simple conditional mean forecast $\hat{w}_{\tau|t}$ for $\tau = t+1,\ldots,t+H$, obtained as

$$\hat{w}_{\tau|t} = (5, 20, 35)P^{t-\tau}s_t, \quad \tau = t+1,\ldots,t+H,$$

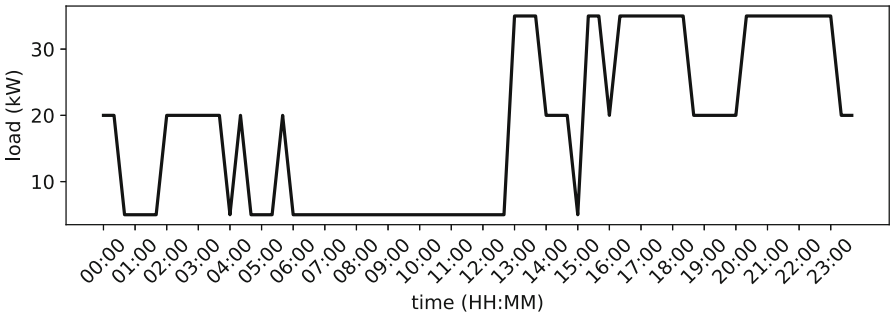

Fig. 8 Simulated load data for January 1st 2020

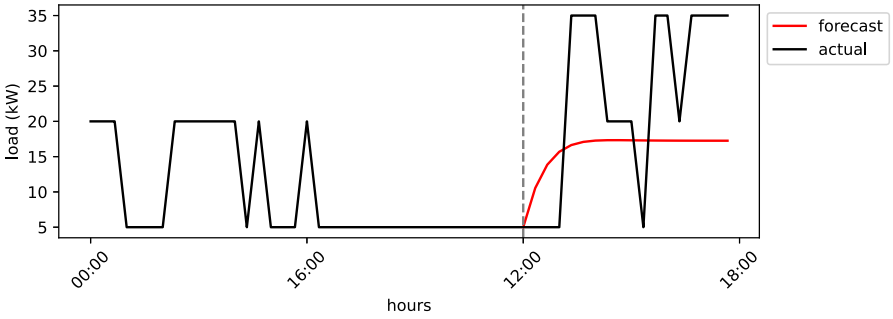

Fig. 9 Actual load and forecasts on January 1st 2020

where s_t is the current state, *i.e.*, $s_t = (1, 0, 0)$ in the low power state, $s_t = (0, 1, 0)$ is the medium power state, and $s_t = (0, 0, 1)$ in the high power state.

We show the load data and 6 hour ahead forecasts for January 1st 2020 in Figure 9. We observe that the forecasts converge to steady state after around 2 hours.

4.5 Simulation results

Using the MPC method in §4.2, data in §4.3 and forecast in §4.4, we change the cost of battery aging γ and simulate in 20 minute intervals (*i.e.*, $\delta = 1/3$ (h)) until the end of life of the battery. Horizon H is set to 6 hours. We observed that the choice of η did not affect the results significantly and we set it to $\eta = 0.5$ for all simulations.

RMSD (*i.e.*, $\mathcal{D}$) of the smoothed load versus battery lifetime is shown in Figure 10. $\mathcal{D} = 10.29$ (kW) for 25 years of simulated load. We observe that with a battery lifetime of 11 years, $\mathcal{D} = 0.44$ (kW) and load is smoothed out approximately 23 times. On the other hand, with a battery lifetime of 22 years, $\mathcal{D} = 6.20$ (kW) and load is smoothed out approximately 1.7 times. Hence $\mathcal{D}$ increases with increasing battery lifetime. This is expected, since with a shorter battery lifetime, the battery is allowed to cycle more aggressively, which results in a smoother signal and lower $\mathcal{D}$. However, this comes at the cost of a shorter battery lifetime.

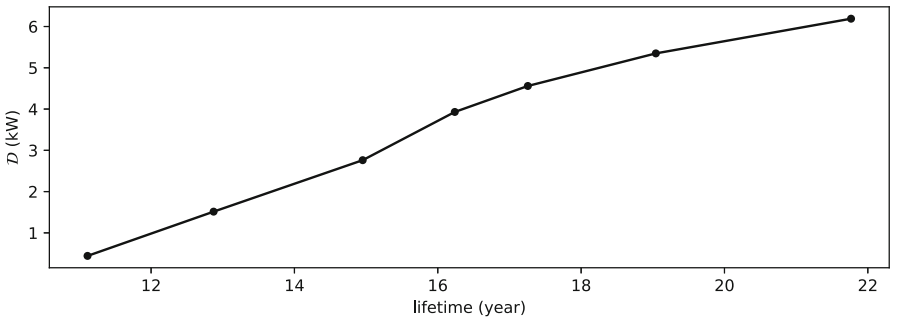

Fig. 10 $\mathcal{D}$ versus battery lifetime

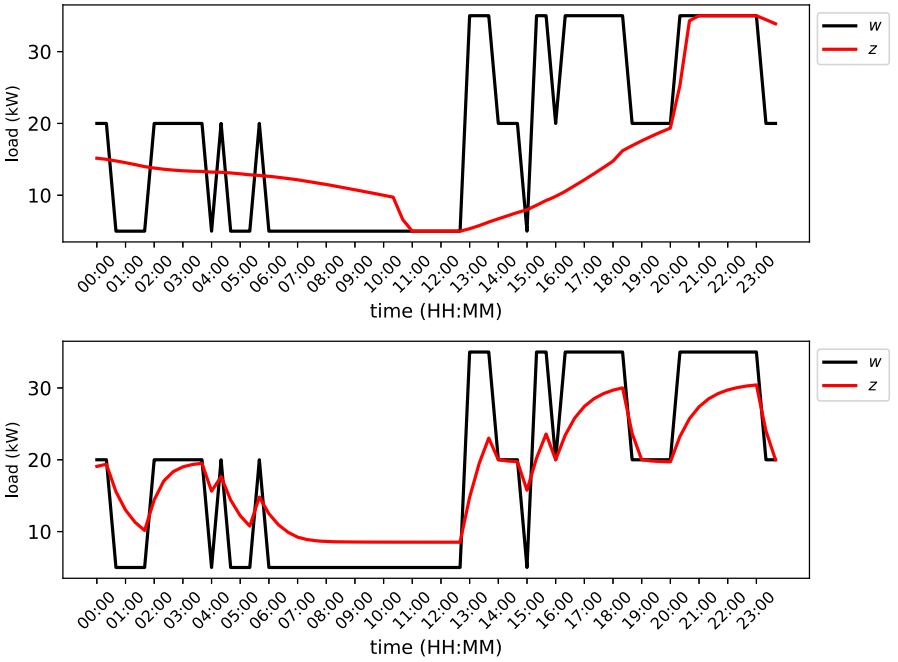

Fig. 11 Daily smoothing profile using forecasts. *Top.* 11 years battery lifetime. *Bottom.* 15 years battery lifetime

We focus on smoothing profiles at a smaller timeframe and look at the daily load smoothing profile for a battery with an 11 year lifetime and 15 year lifetime in Figure 11. The smoothed load changes more slowly than the original load, and as the battery lifetime increases, the smoothing decreases.

5 Conclusion

In this paper we formulated the battery management problem considering two competing objectives: short term goals (with specific examples energy arbitrage and load smoothing) and long term battery lifetime maximization. We adopted an existing

semi-empirical battery aging model for lithium iron phosphate cells and developed a convex optimization-based control strategy using a simplified yet accurate convex approximation of the aging rate.

Our approach employs MPC, leveraging known future prices or load forecasts to optimally balance short term task performance and battery longevity. Through extensive numerical simulations, we clearly demonstrated the trade-off between aggressive short term battery cycling and battery lifetime. We observed that aggressive cycling significantly increased hourly revenues and reduced load fluctuations, but at the cost of a shorter battery lifetime. Conversely, conservative strategies prolonged battery life but yielded lower short term performance gains.

Key novelties of our work include the convex approximation of the battery aging model, enabling computationally efficient optimization within the MPC framework, and systematic quantification of the lifetime-performance trade-off under realistic conditions. Our results underscore the importance of aging-aware optimization strategies, providing a practical and scalable solution that bridges the gap between detailed battery aging models and real-time battery management. Open source implementation and data used in this work are available at ON (2025).

Acknowledgements We would like to thank Prof. Simona Onori for discussions of the battery aging model, Dr. Eric Sager Luxenberg for discussion of an initial formulation, and Maximilian Schaller for careful reading and corrections of the manuscript and code.

Author Contributions The initial idea was proposed by O.N. and S.B. G.O. formulated the convex optimization models, reimplemented and extended the codebase, conducted simulations, and wrote the manuscript. G.O. and O.N. collaborated on discussions related to the aging model. S.B. contributed to the mathematical modeling, supervised experimental results, helped refine notation, improved the clarity of exposition, and edited the paper. P.L. offered general guidance and feedback. All authors reviewed the manuscript.

Funding: Mehmet Giray Ogut's work is supported by the U.S. Department of Energy's Office of Electricity through a grant from the Grid Controls and Communications Division.
Obidike Nnorom Jr. was supported by iBuild Fellowship. Philip Levis was supported by the MemoryDAX Affiliates Program at Stanford University.

Availability of data and materials: All data and code supporting the findings of this study are publicly available at: https://github.com/cvxgrp/aging_aware_battery_control.

Declarations

Ethics approval and consent to participate Not applicable.

Consent for publication Not applicable.

Competing interests The authors declare no competing interests.

References

Abdulla K, de Hoog J, Muenzel V, Suits F, Steer K, Wirth A, Halgamuge S (2018) Optimal operation of energy storage systems considering forecasts and battery degradation. IEEE Transactions on Smart Grid 9(3):2086–2096. https://doi.org/10.1109/tsg.2016.2606490

Abughalieh K, Alawneh S (2019) A survey of parallel implementations for model predictive control. IEEE Access 7:34348–34360

Afram A, Janabi-Sharifi F (2014) Theory and applications of HVAC control systems-a review of model predictive control (MPC). Build Environ 72:343–355
Aksanli B, Rosing T, Pettis E (2013) Distributed battery control for peak power shaving in datacenters. In: Proceedings of the 2013 International Green Computing Conference, pp. 1–8. IEEE
Allam A, Onori S (2020) Online capacity estimation for lithium-ion battery cells via an electrochemical model-based adaptive interconnected observer. IEEE Trans Control Syst Technol 29(4):1636–1651
Barroso L, Hölzle U, Ranganathan P (2019) The Datacenter as a Computer: Designing Warehouse-Scale Machines. Springer Nature, Berlin
Bashir N, Sardar H, Nasir M, Hassan N, Khan H (2017) Lifetime maximization of lead-acid batteries in small scale UPS and distributed generation systems. In: 2017 IEEE Manchester Powertech, pp. 1–6. https://doi.org/10.1109/ptc.2017.7980993. https://ieeexplore.ieee.org/document/7980993/?arnumber=7980993&tag=1
Bellman R (1957) Dynamic Programming. Princeton University Press, Princeton, NJ
Bemporad A, Morari M (2007) Robust model predictive control: A survey. In: Robustness in Identification and Control. Springer, Berlin, pp 207–226
Bertsimas D, Popescu I (2003) Revenue management in a dynamic network environment. Transp Sci 37(3):257–277
Borrelli F, Bemporad A, Morari M (2017) Predictive Control for Linear and Hybrid Systems. Cambridge University Press, Cambridge
Bradbury K, Pratson L, Patiño-Echeverri D (2014) Economic viability of energy storage systems based on price arbitrage potential in real-time us electricity markets. Appl Energy 114:512–519
Camacho E, Bordons C, Camacho E, Bordons C (2007) Constrained Model Predictive Control. Springer, Berlin
Campo P, Morari M (1987) Robust model predictive control. In: 1987 American Control Conference, pp. 1021–1026. IEEE
Carli R, Cavone G, Epicoco N, Scarabaggio P, Dotoli M (2020) Model predictive control to mitigate the Covid-19 outbreak in a multi-region scenario. Annu Rev Control 50:373–393
Chung C, Jangra S, Lai Q, Lin X (2020) Optimization of electric vehicle charging for battery maintenance and degradation management. IEEE Transactions on Transportation Electrification 6(3):958–969. https://doi.org/10.1109/tte.2020.3000181. https://ieeexplore.ieee.org/document/9108305/?arnumber=9108305&tag=1
Collath N, Tepe B, Englberger S, Jossen A, Hesse H (2022) Aging aware operation of lithium-ion battery energy storage systems: A review. Journal of Energy Storage 55:105634. https://doi.org/10.1016/j.est.2022.105634. https://www.sciencedirect.com/science/article/pii/S2352152X2201622X
Critchlow J, Denman A (2017) Embracing the next energy revolution: Electricity storage. Bain & Company, Boston
Cutler C (1979) Dynamic matrix control – a computer control algorithm. In: Proceedings of the Joint Automatic Control Conference
Diamond S, Boyd S (2016) CVXPY: A python-embedded modeling language for convex optimization. J Mach Learn Res 17(83):1–5
Díaz-González F, Sumper A, Gomis-Bellmunt O, Bianchi F (2013) Energy management of flywheel-based energy storage device for wind power smoothing. Appl Energy 110:207–219
Ding Y, Wang L, Li Y, Li D (2018) Model predictive control and its application in agriculture: A review. Comput Electron Agric 151:104–117
El-Naga A, Marei M, El-Goharey H (2017) Second order adaptive notch filter based wind power smoothing using flywheel energy storage system. In: 2017 Nineteenth International Middle East Power Systems Conference (MEPCON), pp. 314–319. IEEE
Elkomy A, Huzayyin A, Abdo T, Adly A, Yassin H (2017) Enhancement of wind energy conversion systems active and reactive power control via flywheel energy storage systems integration. In: 2017 Nineteenth International Middle East Power Systems Conference (MEPCON), pp. 1151–1156. IEEE
ERCOT (2012) Historical dam load zone and hub prices. https://www.ercot.com/mp/data-products/data-product-details?id=np4-180-er
Eren U, Prach A, Koçer B, Raković S, Kayacan E, Açıkmeşe B (2017) Model predictive control in aerospace systems: Current state and opportunities. J Guid Control Dyn 40(7):1541–1566
Felez J, Kim Y, Borrelli F (2019) A model predictive control approach for virtual coupling in railways. IEEE Trans Intell Transp Syst 20(7):2728–2739

Franco A, Doublet M, Bessler W (2016) Physical Multiscale Modeling and Numerical Simulation of Electrochemical Devices for Energy Conversion and Storage: From Theory to Engineering to Practice. Green Energy and Technology. Springer, London. https://doi.org/10.1007/978-1-4471-5677-2

Garcia C, Prett D, Morari M (1989) Model predictive control: Theory and practice — a survey. Automatica 25(3):335–348

Garey M, Johnson D (1990) Computers and Intractability; A Guide to the Theory of NP-completeness. W. H. Freeman & Co., USA

Goulart P, Chen Y (2024) Clarabel: An interior-point solver for conic programs with quadratic objectives. https://doi.org/10.48550/arXiv.2405.12762

Govindan S, Sivasubramaniam A, Urgaonkar B (2011) Benefits and limitations of tapping into stored energy for datacenters. In: Proceedings of the 38th Annual International Symposium on Computer Architecture, pp. 341–352

Govindan, S., Wang, D., Sivasubramaniam, A., Urgaonkar, B.: Leveraging stored energy for handling power emergencies in aggressively provisioned datacenters. In: Proceedings of the Seventeenth International Conference on Architectural Support for Programming Languages and Operating Systems, pp. 75–86 (2012)

Grüne L, Pannek J, Grüne L, Pannek J (2017) Nonlinear Model Predictive Control. Springer, Berlin

Guo Y, Ding Z, Fang Y, Wu D (2011) Cutting down electricity cost in internet data centers by using energy storage. In: 2011 IEEE Global Telecommunications Conference-GLOBECOM 2011, pp. 1–5. IEEE

He G, Ciez R, Moutis P, Kar S, Whitacre J (2020) The economic end of life of electrochemical energy storage. Appl Energy 273:115151

Heirung T, Paulson J, O'Leary J, Mesbah A (2018) Stochastic model predictive control—how does it work? Computers & Chemical Engineering 114:158–170

Holkar K, Waghmare L (2010) An overview of model predictive control. International Journal of Control and Automation 3(4):47–63

Hovgaard T, Boyd S, Jørgensen J (2015) Model predictive control for wind power gradients. Wind Energy 18(6):991–1006

Hu J, Shan Y, Guerrero J, Ioinovici A, Chan K, Rodriguez J (2021) Model predictive control of microgrids - an overview. Renew Sustain Energy Rev 136:110422

Huang Y, Wang H, Khajepour A, He H, Ji J (2017) Model predictive control power management strategies for HEVs: A review. J Power Sources 341:91–106

Jerbi L, Krichen L, Ouali A (2009) A fuzzy logic supervisor for active and reactive power control of a variable speed wind energy conversion system associated to a flywheel storage system. Electric Power Systems Research 79(6):919–925

Jin X, Vora A, Hoshing V, Saha T, Shaver G, Wasynczuk O, Varigonda S (2018) Applicability of available li-ion battery degradation models for system and control algorithm design. Control Eng Pract 71:1–9

Jr. ON, Ogut, G, Boyd S, Levis P (2025) Aging-aware battery control via convex optimization. https://github.com/cvxgrp/aging_aware_battery_control

Keil J, Jossen A (2020) Electrochemical modeling of linear and nonlinear aging of lithium-ion cells. J Electrochem Soc 167(11):110535

Khalid M, Savkin A (2009) Model predictive control for wind power generation smoothing with controlled battery storage. In: Proceedings of the 48th IEEE Conference on Decision and Control (CDC), pp. 7849–7853. IEEE

Khalid M, Savkin A (2010) A model predictive control approach to the problem of wind power smoothing with controlled battery storage. Renewable Energy 35(7):1520–1526

Kontorinis V, Zhang L, Aksanli B, Sampson J, Homayoun H, Pettis E, Tullsen D, Rosing T (2012) Managing distributed UPS energy for effective power capping in data centers. ACM Sigarch Computer Architecture News 40(3):488–499

Latz A, Zausch J (2013) Thermodynamic derivation of a butler-volmer model for intercalation in li-ion batteries. Electrochim Acta 110:358–362

Lazic N, Boutilier C, Lu T, Wong E, Roy B, Ryu M, Imwalle G (2018) Data center cooling using model-predictive control. Adv Neural Inf Process Syst **31**

Li J, Wang D, Deng L, Cui Z, Lyu C, Wang L, Pecht M (2020) Aging modes analysis and physical parameter identification based on a simplified electrochemical model for lithium-ion batteries. Journal of Energy Storage 31:101538

Li Y, Mughees M, Chen Y, Li Y (2024) The unseen AI disruptions for power grids: LLM-induced transients. arXiv Preprint arXiv:2409.11416

Liu K, Ashwin T, Hu X, Lucu M, Widanage W (2020) An evaluation study of different modelling techniques for calendar ageing prediction of lithium-ion batteries. Renew Sustain Energy Rev 131:110017

Liu K, Zou C, Li K, Wik T (2018) Charging pattern optimization for lithium-ion batteries with an electrothermal-aging model. IEEE Trans Industr Inf 14(12):5463–5474. https://doi.org/10.1109/tii.2018.2866493. https://ieeexplore.ieee.org/document/8444057/?arnumber=8444057

Maheshwari A, Paterakis N, Santarelli M, Gibescu M (2020) Optimizing the operation of energy storage using a non-linear lithium-ion battery degradation model. Appl Energy 261:114360. https://doi.org/10.1016/j.apenergy.2019.114360. https://www.sciencedirect.com/science/article/pii/s0306261919320471

Mamun A, Narayanan I, Wang D, Sivasubramaniam A, Fathy H (2016) Multi-objective optimization of demand response in a datacenter with lithium-ion battery storage. Journal of Energy Storage 7:258–269. https://doi.org/10.1016/j.est.2016.08.002. https://www.sciencedirect.com/science/article/pii/s2352152x16301025

Mamun A, Wang D, Narayanan I, Sivasubramaniam A, Fathy H (2015) Physics-based simulation of the impact of demand response on lead-acid emergency power availability in a datacenter. J Power Sources 275:516–524

Marano V, Onori S, Guezennec Y, Rizzoni G, Madella N (2009) Lithium-ion batteries life estimation for plug-in hybrid electric vehicles. In: 2009 IEEE Vehicle Power and Propulsion Conference, pp. 536–543. IEEE

Mayne D (2014) Model predictive control: Recent developments and future promise. Automatica 50(12):2967–2986

Mayne D, Seron M, Raković S (2005) Robust model predictive control of constrained linear systems with bounded disturbances. Automatica 41(2):219–224

McConnell D, Forcey T, Sandiford M (2015) Estimating the value of electricity storage in an energy-only wholesale market. Appl Energy 159:422–432

Mesbah A (2016) Stochastic model predictive control: An overview and perspectives for future research. IEEE Control Syst Mag 36(6):30–44

Metz D, Saraiva J (2018) Use of battery storage systems for price arbitrage operations in the 15-and 60-min German intraday markets. Electric Power Systems Research 160:27–36

Miller C, Goutham M, Chen X, Hanumalagutti P, Blaser R, Stockar S (2022) A semi-empirical approach to a physically based aging model for home energy management systems. In: 2022 IEEE Conference on Control Technology and Applications (CCTA), pp. 165–170. IEEE

Pelletier S, Jabali O, Laporte G, Veneroni M (2017) Battery degradation and behaviour for electric vehicles: Review and numerical analyses of several models. Transportation Research Part B Methodological 103:158–187

Péni T, Csutak B, Szederkényi G, Röst G (2020) Nonlinear model predictive control with logic constraints for Covid-19 management. Nonlinear Dyn 102:1965–1986

Primbs J (2009) Dynamic hedging of basket options under proportional transaction costs using receding horizon control. Int J Control 82(10):1841–1855

Raimondo D, Limon D, Lazar M, Magni L, Camacho E (2009) Min-max model predictive control of nonlinear systems: A unifying overview on stability. Eur J Control 15(1):5–21

Rakovic S, Levine W (2018) Handbook of Model Predictive Control. Springer, Berlin

Rawlings J, Mayne D, Diehl M (2017) Model Predictive Control: Theory, Computation, and Design, vol 2. Nob Hill Publishing Madison, WI

Ren C, Wang D, Urgaonkar B, Sivasubramaniam A (2012) Carbon-aware energy capacity planning for datacenters. In: Proceedings of the 2012 IEEE 20th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems, pp. 391–400. IEEE

Schmidt O, Hawkes A, Gambhir A, Staffell I (2017) The future cost of electrical energy storage based on experience rates. Nat Energy 2(8):1–8

Schwenzer M, Ay M, Bergs T, Abel D (2021) Review on model predictive control: An engineering perspective. The International Journal of Advanced Manufacturing Technology 117(5):1327–1349

Serrao L, Onori S, Rizzoni G, Guezennec Y (2009) A novel model-based algorithm for battery prognosis. IFAC Proceedings Volumes 42(8):923–928. https://doi.org/10.3182/20090630-4-es-2003.00152. https://www.sciencedirect.com/science/article/pii/s1474667016358955

Shen X, Boyd S (2021) Incremental proximal multi-forecast model predictive control. arXiv Preprint arXiv:2111.14728

Sioshansi R, Denholm P, Jenkin T, Weiss J (2009) Estimating the value of electricity storage in PJM: Arbitrage and some welfare effects. Energy Economics 31(2):269–277

Staffell I, Rustomji M (2016) Maximising the value of electricity storage. Journal of Energy Storage 8:212–225

Suri G, Onori S (2016) A control-oriented cycle-life model for hybrid electric vehicle lithium-ion batteries. Energy 96:644–653. https://doi.org/10.1016/j.energy.2015.11.075. https://www.sciencedirect.com/science/article/pii/s0360544215016382

Suvire G, Molina M, Mercado P (2012) Improving the integration of wind power generation into AC microgrids using flywheel energy storage. IEEE Transactions on Smart Grid 3(4):1945–1954

Talluri K, Ryzin GV (2006) The Theory and Practice of Revenue Management, vol 68. Springer Science & Business Media, Berlin

Torregrosa A, Broatch A, Olmeda P, Agizza L (2024) A semi-empirical model of the calendar ageing of lithium-ion batteries aimed at automotive and deep-space applications. Journal of Energy Storage 80:110388

Urgaonkar R, Urgaonkar B, Neely M, Sivasubramaniam A (2011) Optimal power cost management using stored energy in data centers. In: Proceedings of the ACM Sigmetrics Joint International Conference on Measurement and Modeling of Computer Systems, pp. 221–232

Vermeer W, Mouli G, Bauer P (2021) A comprehensive review on the characteristics and modeling of lithium-ion battery aging. IEEE Transactions on Transportation Electrification 8(2):2205–2232

Wang D, Govindan S, Sivasubramaniam A, Kansal A, Liu J, Khessib B (2014) Underprovisioning backup power infrastructure for datacenters. In: Proceedings of the 19th International Conference on Architectural Support for Programming Languages and Operating Systems, pp. 177–192

Wang D, Ren C, Sivasubramaniam A, Urgaonkar B, Fathy H (2012) Energy storage in datacenters: What, where, and how much? In: Proceedings of the 12th ACM Sigmetrics/performance Joint International Conference on Measurement and Modeling of Computer Systems, pp. 187–198

Wang Y, Boyd S (2009) Fast model predictive control using online optimization. IEEE Trans Control Syst Technol 18(2):267–278

Weitzel T, Glock C (2018) Energy management for stationary electric energy storage systems: A systematic literature review. Eur J Oper Res 264(2):582–606. https://doi.org/10.1016/j.ejor.2017.06.052. https://www.sciencedirect.com/science/article/pii/s0377221717305933

Wilson I, Barbour E, Ketelaer T, Kuckshinrichs W (2018) An analysis of storage revenues from the time-shifting of electrical energy in Germany and Great Britain from 2010 to 2016. Journal of Energy Storage 17:446–456

Xiong R, Pan Y, Shen W, Li H, Sun F (2020) Lithium-ion battery aging mechanisms and diagnosis method for automotive applications: Recent advances and perspectives. Renew Sustain Energy Rev 131:110048

Xu B, Zhao J, Zheng T, Litvinov E, Kirschen D (2017) Factoring the cycle aging cost of batteries participating in electricity markets. IEEE Trans Power Syst 33(2):2248–2259

Xu B, Zhao J, Zheng T, Litvinov E, Kirschen D (2018) Factoring the cycle aging cost of batteries participating in electricity markets. IEEE Trans Power Syst 33(2):2248–2259. https://doi.org/10.1109/tpwrs.2017.2733339

Zafirakis D, Chalvatzis K, Baiocchi G, Daskalakis G (2016) The value of arbitrage for energy storage: Evidence from European electricity markets. Appl Energy 184:971–986

Zhang P, Yan F, Du C (2015) A comprehensive analysis of energy management strategies for hybrid electric vehicles based on bibliometrics. Renew Sustain Energy Rev 48:88–104