Degradation-Model-Agnostic Energy Throughput Control for EV Range Preservation and Battery Lifetime Extension via Cell-Level Inverters

Citation for the original published paper (version of record):
Jiang, S., Tao, S., Molina, V. et al (2026). Degradation-Model-Agnostic Energy Throughput Control for EV Range Preservation and Battery
Lifetime Extension via Cell-Level Inverters. IEEE Transactions on Transportation Electrification, In Press. http://dx.doi.org/10.1109/TTE.2026.3725087

Degradation-Model-Agnostic Energy Throughput Control for EV Range Preservation and Battery Lifetime Extension via Cell-Level Inverters

Shida Jiang, Shengyu Tao, Vincent Molina, Junzhe Shi, and Scott Moura

Abstract*—A conventional electric vehicle (EV) powertrain relies on a centralized high-voltage DC–AC inverter, thereby limiting cell-level control and potentially reducing overall driving range and battery lifetime. This paper studies an H-bridge-based cell-level inverter topology that performs power conversion at the cell level, enabling independent control of individual cells and expanding the design space for battery management. Leveraging these additional degrees of freedom, we propose a degradation-model-agnostic energy-throughput control strategy that preserves usable energy and EV range while extending battery-pack lifetime. Because usable energy (and thus driving range) and lifetime are governed by the cells with the lowest state-of-charge (SOC) and state-of-health (SOH), respectively, the proposed controller preferentially routes energy throughput to healthier cells. Specifically, during charging, it permits cell SOCs to diverge to promote SOH equalization; during discharging, it balances remaining usable capacity to maximize usable energy under per-cell constraints. The proposed SOC–SOH-aware control strategy is evaluated on two aging models representing lithium manganese oxide (LMO) and lithium iron phosphate (LFP) chemistries, using a Tesla Model 3 charge–discharge profile under a broad set of aging-model, controller-design, end-of-life (EOL), and estimation assumptions. In the reference case, which includes zero-mean SOH-estimation noise with a standard deviation of 2 percentage points, simulations show lifetime improvements of 13.0 % and 8.9 % for LFP and LMO batteries, respectively. These results show that the proposed controller can extend battery-pack lifetime without reducing usable energy, while requiring neither a specific aging model nor prior knowledge of the next trip. The code and supplementary material for this study are available at https://github.com/Shida-Jiang/Model-Agnostic-Control-for-EV-Range-and-Lifetime-Extension.

***Index Terms*—Battery management system, State-of-charge, State-of-health, Electric vehicle, Cascaded H-bridge, Level-shifted pulse-width modulation**

I. Introduction

Electric vehicles (EVs) rely on power electronic converters to transform the battery pack's DC output into the three-phase AC supply required by the traction motor. As shown in Fig. 1(a), in conventional EV powertrains, this conversion is performed by a centralized high-voltage inverter, and the battery pack is operated largely as a single series-connected energy source. As a result, usable capacity is constrained by the weakest cell in the string, and opportunities for cell-level energy management are limited.

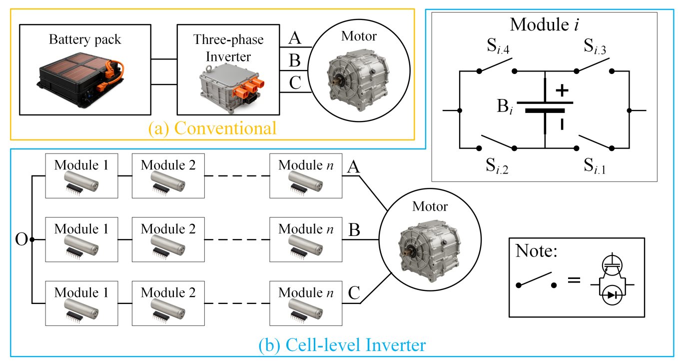

Fig. 1: Comparison between the (a) conventional inverter powertrain and (b) the H-bridge-based cell-level inverter powertrain. The new topology introduces new degrees of freedom for range and lifetime extension.

Cell-level inverter powertrains, such as the cascaded H-bridge shown in Fig. 1(b), provide a fundamentally different architecture by performing power conversion at the cell or module level [1], [2]. This topology introduces additional degrees of freedom: individual cells can be activated, bypassed, charged, or discharged at different effective rates while still producing the commanded motor voltage. Existing work has mainly used these degrees of freedom for state-of-charge (SOC) balancing, fault tolerance, harmonic reduction, or charging control [1]–[3]. Among these usages, SOC balancing is useful for maximizing extractable energy and hence EV range. Tresca et al., for example, experimentally tested a reconfigurable cascaded multilevel converter that simultaneously performed power conversion and SOC balancing for up to nine battery cells using a single controller [4], [5]. Recently, Stellantis and Saft began real-world testing of an EV prototype based on a similar cell-level inverter topology, demonstrating its potential to improve efficiency and extend battery lifetime [6], [7].

However, SOC balancing alone does not directly address battery-pack lifetime. Lithium-ion cells can age at different rates even under identical usage, and thermal or manufacturing heterogeneity can further amplify cell-to-cell state-of-health (SOH) dispersion [8], [9]. Since pack lifetime is ultimately limited by the weakest or retired cells, lifetime can, in

Shida Jiang, Junzhe Shi, and Scott Moura are with the Department of Civil and Environmental Engineering, University of California, Berkeley, Berkeley, CA 94720 USA. (email: shida_jiang@berkeley.edu, junzhe@berkeley.edu, smoura@berkeley.edu)

Shengyu Tao is with the Department of Civil and Environmental Engineering, University of California, Berkeley, Berkeley, CA 94720 USA, and also with the Department of Electrical Engineering, Chalmers University of Technology, 41296 Gothenburg, Sweden. (email: shengyu.tao@chalmers.se)

Vincent Molina was with the BMW Group Technology Office USA, Mountain View, CA 94043 USA (e-mail: molina.vincent@hotmail.fr).

This work was funded by BMW USA. (Corresponding author: Shengyu Tao)

This article has been accepted for publication in IEEE Transactions on Transportation Electrification. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TTE.2026.3725087

principle, be extended by routing more energy throughput to healthier cells. The difficulty is that SOH balancing conflicts with SOC balancing: aggressively using healthier cells can create SOC imbalance and leave usable energy stranded at the end of discharge, reducing EV range.

Because of this apparent contradiction, very few prior studies have considered SOH in cell-balancing control. Xia and Abu Qahouq incorporate estimated cell capacity into an SOC-balancing controller, while Kim et al. compare SOC-based and terminal-voltage-based balancing and quantify their effects on battery lifetime [10], [11]. In both studies, however, the control objective remains SOC equalization or improved energy utilization, rather than deliberate redistribution of cell throughput to equalize long-term aging. Pröbstl et al. directly target lifetime extension by reducing the load on low-SOH cells [12]. However, their controller assumes that the required charge and duration of the upcoming trip are known in advance. Yu et al. also consider joint SOC–SOH balancing, but for distributed multi-battery systems connected to a common busbar [13]. Their method allocates prescribed power commands among battery units and explicitly relies on a semi-empirical aging model. It therefore addresses a different scheduling problem and does not remove dependence on power-demand information or a specific aging model.

The prior studies therefore leave an important gap. A practical lifetime-oriented controller must intentionally redistribute cell throughput to equalize long-term aging, while ensuring that the resulting SOC imbalance does not reduce the pack's usable capacity during an unknown future trip. It should also avoid dependence on a specific battery aging model, whose structure and parameters are difficult to identify reliably. To the best of our knowledge, no existing method satisfies these requirements simultaneously.

To address this gap, we propose a strategy that prolongs battery life without sacrificing usable capacity (and thus EV range). During charging, the controller deliberately allows cell SOCs to diverge so that healthier cells receive more throughput. During discharge, it balances the remaining usable capacity so that the charging imbalance does not strand energy in the pack. However, despite the appeal of this idea, a direct end-to-end formulation is difficult. Guaranteeing usable capacity during the subsequent trip would ordinarily require its discharge trajectory. In addition, optimizing cell-level operation throughout charging requires time-indexed discrete cell-to-level assignments together with continuous current and voltage commands. The resulting problem is a large-scale mixed-integer optimization whose size grows with the charging horizon.

We overcome these difficulties through a theoretically grounded reformulation. The achievability result in Section III-B allows the controller to optimize target cell charge gains rather than the complete sequence of switching assignments. A discharge-feasibility constraint preserves subsequent usable capacity without requiring a forecast of the drive trajectory. The online optimization contains no degradation law or identified aging parameters. Instead, it uses battery-state and operating information normally available to the battery management system (BMS), including cell SOC, SOH, or usable-capacity estimates, the open-circuit voltage (OCV) function, representative impedance, and current and voltage limits. Its tunable SOH-based weighting law protects weaker cells by directing more throughput to healthier cells, without predicting their degradation rates.

The main contributions of this paper are:

- We propose a control strategy that reconciles SOC and SOH balancing for battery lifetime extension while requiring neither a predicted trip nor an explicit aging model.
- We derive a necessary and sufficient achievability condition for cell-level throughput allocation under level-shifted pulse-width modulation (LSPWM). This result converts a time-indexed mixed-integer problem into a tractable linear program over stage-wise cell charge gains.
- We evaluate the strategy using lithium iron phosphate (LFP) and lithium manganese oxide (LMO) aging models, real charging records from a Tesla Model 3, and a broad set of aging-model, objective-function, end-of-life (EOL), and estimation assumptions. The reference-case lifetime improvements are 13.0 % and 8.9 % for LFP and LMO batteries, respectively.

The remainder of the paper is organized as follows. Section 2 introduces key definitions related to the battery pack's health and remaining capacity. Section 3 introduces the operation principle of the cell-level inverter topology and the feasible region of its control strategy. Section 4 presents our control strategy. Section 5 presents the simulation study. Finally, the conclusions are drawn in Section 6.

II. Key Definitions

A. Cell State-of-charge and State-of-health

In this paper, the remaining capacity and health of cells are defined by two states: SOC and SOH. Throughout the equations and simulations, SOC and SOH are both represented as dimensionless fractions between 0 and 1. As described in the introduction, SOC describes the normalized remaining capacity of a cell. Namely,

$$SOC = \frac{Q_r}{Q_{max}}, \tag{1}$$

where Q_r is the remaining capacity of the cell, and Q_{max} is the present maximum capacity of the cell. SOH describes the normalized present maximum capacity of a cell. According to [14],

$$SOH = \frac{Q_{max}}{Q_{no}}, \tag{2}$$

where Q_{no} is the nominal capacity of the cell.

B. End-of-life, Pack SOC, and Pack SOH

Battery EOL is commonly associated with a cell SOH threshold of approximately 70–80 % for EV applications [15]. However, the appropriate threshold is application-dependent: some retired EV cells still satisfy daily driving requirements below 70 % SOH [16], [17], while regulatory and testing documents often relate EOL to usable capacity, power capability, or retained driving range rather than a universal cell-level SOH threshold [18]–[20]. For the cell-level inverter topology

This article has been accepted for publication in IEEE Transactions on Transportation Electrification. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TTE.2026.3725087

considered here, we distinguish cell EOL from pack EOL. A cell reaches EOL and is bypassed when its SOH falls below the cell-level threshold $SOH_{\text{cell,EOL}}$. The pack reaches EOL when its usable pack SOH falls below the pack-level threshold $SOH_{\text{pack,EOL}}$. Unless otherwise stated, $SOH_{\text{cell,EOL}} = 60\,\%$ and $SOH_{\text{pack,EOL}} = 70\,\%$. Alternative combinations of these thresholds are evaluated in Section V-C.

Following [21], pack SOH is defined as

$$SOH_{\text{pack}} = \frac{Q_{\text{max,pack}}}{Q_{\text{no,pack}}}, \tag{3}$$

where $Q_{\text{max,pack}}$ is the maximum usable discharge capacity and $Q_{\text{no,pack}}$ is the nominal pack capacity. This definition assumes sufficient voltage and power headroom for all active cells to reach their respective lower cutoff limits during discharge.

For the conventional topology in Fig. 1(a), SOH_{pack} equals the minimum SOH among all cells, since usable pack capacity is limited by the lowest-capacity cell [22]. For the cell-level inverter topology in Fig. 1(b), the usable pack capacity equals the sum of the present maximum usable capacities of cells that have not reached cell EOL. When a cell reaches $SOH_{\text{cell,EOL}}$, it is bypassed and no longer contributes to $Q_{\text{max,pack}}$. However, the pack can remain in service until SOH_{pack} reaches $SOH_{\text{pack,EOL}}$. Thus, the EOL of an individual cell is distinct from the EOL of the complete pack, although the latter is still affected by the former.

Meanwhile, the pack SOC is defined as

$$SOC_{\text{pack}} = \frac{Q_{r,\text{pack}}}{Q_{\text{max,pack}}}, \tag{4}$$

where $Q_{r,\text{pack}}$ is the remaining usable capacity of the pack.

III. Operating Principle and Control of the Cell-Level Inverter Topology

A. Level-shifted Pulse Width Modulation

Existing literature has proposed many different types of control methods for the cascaded H-bridge topology shown in Figure 1(b). The most popular ones include phase-shifted pulse width modulation (PSPWM), LSPWM, selective harmonics elimination (SHE), and space-vector pulse width modulation (SVPWM) [23], [24]. However, most of these control techniques are designed for balanced DC sources, meaning that each H-bridge is used equally in the long run. Here, however, we want to charge and discharge different cells at different rates so that their SOH can be more balanced. Therefore, the control technique we are especially interested in is LSPWM, and its key idea is illustrated in Figure 2.

As shown in Fig. 2(a), the modulation wave (desired phase voltage reference) is quantized into several discrete *voltage levels*. Each cell (H-bridge) is assigned to a voltage level, and the quantization step is approximately the cell/module DC voltage. For illustration, we assume all cell voltages are equal to U_{dc}, although in practice they may differ slightly.

After the level assignment, each H-bridge generates a PWM waveform that tracks the duty cycle associated with its assigned level, as shown in Fig. 2(b). The modulation wave can be sinusoidal or quasi-constant, depending on the charging/discharging scenario. For AC operation, a sinusoidal modulation wave is used; cells assigned to higher voltage

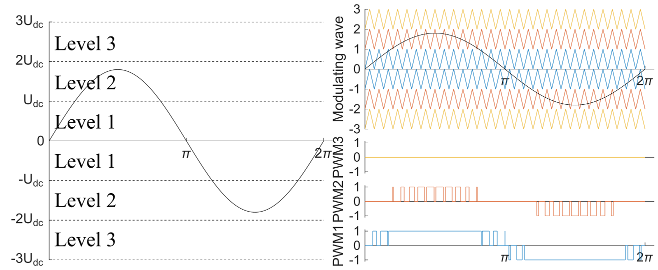

(a) Voltage level division of the modulation wave

(b) PWM wave generation

Fig. 2: The basic idea of level-shifted pulse width modulation control

levels are activated less frequently than those assigned to lower levels. For DC charging, the modulation wave is instead a slowly varying quasi-constant voltage. In this case, almost all voltage levels have duty cycles equal to either zero or one, and at most one voltage level has a non-binary duty cycle. The level assignment can be updated frequently, enabling either nearly equal long-term usage or intentionally unbalanced usage to support SOH balancing.

B. Imbalanced operation within each phase and related theoretical analysis

While prior literature also noted LSPWM's performance under an imbalanced power distribution [25], [26], these works typically treated imbalance as a non-ideal condition and rarely deliberately exploited it. Here, we embrace and leverage LSPWM's ability to distribute throughput unevenly across cells. This raises a natural question: what is the maximum imbalance that LSPWM can realize? Equivalently, given the initial cell states, which final states are achievable under LSPWM?

Indeed, although LSPWM enables different cells to be charged at different rates, the achievable imbalance is still constrained by the duty-cycle structure of the modulation reference. In the illustrative case of Fig. 2, the first two voltage levels have nonzero duty cycles, meaning that at least two H-bridges must be active during each modulation period; thus, it is impossible to charge only one cell while charging none of the others. On the other hand, it is possible to charge Cells 1 and 2 equally while bypassing Cell 3 on average. For example, assign Cell 1 to Level 1, Cell 2 to Level 2, and Cell 3 to Level 3 during one modulation period, and swap the assignments of Cells 1 and 2 in the next period. Repeating this swapping yields equal average throughput for Cells 1 and 2 while Cell 3 remains bypassed. In the general case, Theorem 1 provides a necessary and sufficient condition relating the initial and final states. We first introduce the following definitions.

Consider one phase leg comprising n_1 series-connected H-bridge cells. Let $Q_{\text{initial},i}$ and $Q_{\text{final},i}$ denote the initial charge capacity and the desired final charge capacity of cell i, respectively, and define $\Delta Q_i = Q_{\text{final},i} - Q_{\text{initial},i}$. Let d_ℓ denote the duty cycle of voltage level ℓ in Fig. 2(b). We assume that the controller may reassign cells to voltage levels from one modulation period to the next.

This article has been accepted for publication in IEEE Transactions on Transportation Electrification. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TTE.2026.3725087

Definition 1 (Achievable)**.** *A target final state* $\{Q_{\mathit{final},i}\}_{i=1}^{n_1}$ *is* achievable *from an initial state* $\{Q_{\mathit{initial},i}\}_{i=1}^{n_1}$ *if there exist a time horizon* $\Delta t > 0$ *and an admissible level-assignment policy such that, under LSPWM operation, all cells reach* $Q_{\mathit{final},i}$ *after* Δt.

Theorem 1. *Assume that (i) all cell DC voltages are equal and constant over time, so the duty-cycle pattern* $\{d_\ell\}_{\ell=1}^{n_1}$ *induced by the modulation reference is constant, and (ii) the modulation period* T_m *is much shorter than the total charging time, i.e.,* $T_m \ll \Delta t$, *so that level assignments can be permuted many times over* Δt. *Let* $\Delta Q_{(1)} \geq \Delta Q_{(2)} \geq \cdots \geq \Delta Q_{(n_1)}$ *be the sorted capacity gains across cells, and let* $d_1 \geq d_2 \geq \cdots \geq d_{n_1}$ *denote the duty cycles across voltage levels (largest to smallest). Then the target final state is achievable if and only if*

$$\frac{\sum_{j=1}^{k} \Delta Q_{(j)}}{\sum_{j=1}^{n_1} \Delta Q_{(j)}} \leq \frac{\sum_{j=1}^{k} d_j}{\sum_{j=1}^{n_1} d_j}, \qquad k = 1, 2, \ldots, n_1, \tag{5}$$

(with equality at $k = n_1$*). Moreover, in LSPWM,* d_i *equals the fraction of time the reference exceeds the* i*th carrier threshold (level boundary); hence* $d_1 \geq d_2 \geq \cdots \geq d_{n_1}$ *for any reference waveform.*

Proof. Let

$$x_i = \frac{\Delta Q_i}{\sum_{r=1}^{n_1} \Delta Q_r}, \qquad p_i = \frac{d_i}{\sum_{r=1}^{n_1} d_r},$$

and sort both vectors in nonincreasing order. During any modulation period, assigning cells to voltage levels applies a permutation of the normalized duty vector p. Over many modulation periods, any admissible policy therefore produces an average allocation x that lies in the convex hull of all permutations of p.

The convex hull of all permutations of p is characterized by the majorization inequalities

$$\sum_{j=1}^{k} x_{(j)} \leq \sum_{j=1}^{k} p_j, \qquad k = 1, \ldots, n_1,$$

with equality at $k = n_1$. Substituting the definitions of x_i and p_i gives (5), proving necessity.

Conversely, if (5) holds, then the desired normalized allocation x is majorized by p, and hence belongs to the convex hull of the permutations of p. Therefore, x can be written as a convex combination of permuted duty-cycle vectors. Implementing these permutations for the corresponding fractions of the charging horizon yields the required average capacity gains $\{\Delta Q_i\}$. Since $T_m \ll \Delta t$, this time-sharing construction can be implemented arbitrarily closely by LSPWM level reassignment. Thus the target final state is achievable. □

Remark. Condition (5) characterizes achievable final states only when the duty cycles $\{d_\ell\}$ are constant over time. In practice, the duty cycles depend on the cells' terminal voltage and the amplitude of the phase voltage, both of which can vary as the charging progresses. Nevertheless, while Theorem 1 does not guarantee achievability under general voltage variations, the charging strategy we mentioned in the sufficiency proof of Theorem 1 is still very useful, as we explain below.

Let $Q_i(t)$ denote the charge content of cell i at time t, and define the remaining target gain $\Delta Q_i^r(t) \triangleq \left| Q_{\mathrm{final},i} - Q_i(t) \right|$, where $Q_{\mathrm{final},i}$ is the target charge content of the ongoing operation ($Q_{\mathrm{final},i} = 0$ when the target is complete depletion). We propose the following strategy.

Strategy 1 (remaining capacity gain balancing). At each modulation period, assign the cell with the jth largest $\Delta Q_i^r(t)$ to voltage level j. If ties occur, preserve them by cyclically rotating the level assignments among the tied cells.

Corollary 1. *Under the assumptions of Theorem 1, Strategy 1 reaches every target satisfying* (5)*, up to an arbitrarily small implementation error as* $T_m/\Delta t \to 0$.

Proof. Normalize the target gains and duty cycles as

$$q_i = \frac{\Delta Q_i}{\sum_{r=1}^{n_1} \Delta Q_r}, \qquad p_i = \frac{d_i}{\sum_{r=1}^{n_1} d_r},$$

and let $\tau \in [0, 1]$ denote normalized charging time. Let $r_{(1)}(\tau) \geq \cdots \geq r_{(n_1)}(\tau)$ denote the sorted remaining normalized gains.

Under Strategy 1, tied cells cyclically share the corresponding voltage levels, so equal residuals remain equal. Consequently, the sorted residuals form contiguous blocks that can merge but cannot split. For any boundary a between two such blocks, the bottom $n_1 - a + 1$ residuals are served by the bottom $n_1 - a + 1$ duty shares. Therefore,

$$\sum_{i=a}^{n_1} r_{(i)}(\tau) = \sum_{i=a}^{n_1} q_{(i)} - \tau \sum_{i=a}^{n_1} p_i.$$

Condition (5), together with equality of the total sums, is equivalent to

$$\sum_{i=a}^{n_1} q_{(i)} \geq \sum_{i=a}^{n_1} p_i.$$

Hence, no block with a positive associated duty share can reach zero before $\tau = 1$. Since

$$\sum_{i=1}^{n_1} r_i(\tau) = 1 - \tau,$$

all residuals reach zero at $\tau = 1$. Thus, Strategy 1 reaches the target. □

Strategy 1 is a constructive realization of the achievability result in Theorem 1. As established in Corollary 1, when the duty-cycle vector is constant and level reassignment is sufficiently frequent, the strategy reaches every achievable target by continually assigning the largest duty shares to the largest remaining target gains. When terminal voltages vary, the duty-cycle vector is no longer constant, so exact attainment is not guaranteed by the corollary. For the cell parameters in Section V-A, the voltage perturbation is nevertheless modest. Namely, at the maximum reference C-rate of about 2.44 C, the modeled ohmic drop is 56 mV per cell for $R_0 = 0.01\,\Omega$ and 281 mV for the conservative $R_0 = 0.05\,\Omega$, i.e., roughly 2 % and 9 % of the cell voltage. Meanwhile, the slower OCV drift is captured by the stage-wise discretization in Section IV. The same feedback rule then continuously adapts the assignments to the remaining target gains, and the conservative-resistance case in Section V-C leaves the lifetime results unchanged.

Strategy 2 (remaining-capacity balancing). At each modulation period, if the battery pack is discharging, assign the cell with the jth largest $\Delta Q_i^r(t)$ to voltage level j; if the battery pack is charging (e.g., in regenerative mode), assign the cell

This article has been accepted for publication in IEEE Transactions on Transportation Electrification. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TTE.2026.3725087

with the jth smallest $\Delta Q_i^r(t)$ to voltage level j. If ties occur, maintain ties by cyclically rotating level assignments among tied cells.

Remark. Strategy 2 can be regarded as a special case of Strategy 1 in which the desired final state is complete depletion of the usable pack capacity. Under the constant-terminal-voltage approximation, the strategy allocates higher Ah throughput to cells with larger remaining usable capacity during discharge, thereby minimizing residual unusable capacity at pack depletion and maximizing extractable energy (and thus range). In the general case, it remains an effective and computationally light heuristic, consistent with the intuition that lower-capacity cells should be discharged more slowly to avoid prematurely limiting the pack.

Strategy 2 provides an effective and computationally efficient way to increase usable discharge capacity and reduce remaining-capacity imbalance. It is closely related to SOC balancing, which is presented below.

Strategy 3 (SOC balancing). At each modulation period, if the battery pack is discharging, assign the cell with the jth largest SOC to voltage level j; if the battery pack is charging (e.g., in regenerative mode), assign the cell with the jth smallest SOC to voltage level j. If ties occur, maintain ties by cyclically rotating level assignments among tied cells.

Strategies 2 and 3 coincide when cells have identical SOH, but they differ when cells exhibit SOH heterogeneity: SOC balancing equalizes SOC, whereas remaining-capacity balancing accounts for differences in usable capacities and can therefore extract more energy under heterogeneous SOH. In practice, SOC balancing is often easier to implement because SOC can be estimated without explicitly estimating SOH (or capacity), whereas remaining-capacity balancing is most accurate when both SOC and SOH are available. Since most existing works do not target SOH balancing, SOC balancing is more commonly adopted than remaining-capacity balancing.

In principle, the three-phase topology also permits power imbalance across phases through common-mode voltage injection: changing the neutral-point voltage modifies phase-voltage magnitudes without changing the commanded line-to-line voltages. In this paper, however, we intentionally exploit imbalance only within each phase. The reason is that within-phase imbalance can be implemented by LSPWM level reassignment without changing the motor-side voltage command, whereas phase-level power imbalance increases RMS phase currents and conduction losses for the same delivered power. For example, if balanced operation gives total conduction loss $3I^2R$, an extreme imbalance with two phases carrying $\sqrt{3}I$ and the third phase carrying zero current gives $6I^2R$. Therefore, phase-level imbalance is not used in the proposed strategy.

IV. Control Strategy

A. Basic Control Idea

The idea of jointly optimizing battery lifetime and energy utilization can be illustrated using a simple example. In Fig. 3, we consider a pack with two cells. Cell 2 has a higher SOH

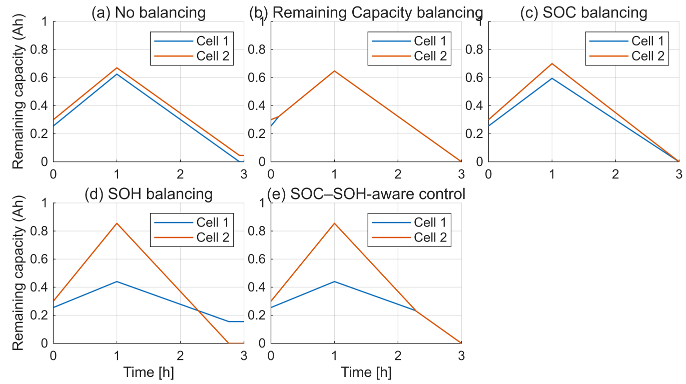

Fig. 3: Comparison of different control methods for the cell-level inverter with two cells in total. The battery pack is first charged from 40 % SOC to 70 % SOC and is then fully discharged.

than Cell 1, and both cells start at 30 % SOC with nominal capacity 1 Ah. The pack is first charged with a fixed current profile (i.e., fixed AC current magnitude on the source side) to a target pack SOC of 70 %, and is then fully discharged. Although the overall current profile is the same in Fig. 3(a)–(e), different strategies distribute the Ah throughput differently across the two cells. If the two cells are charged and discharged equally, the remaining capacities follow Fig. 3(a). Because Cell 1 has lower SOH (and thus lower usable capacity), it is depleted first during discharge, leaving a portion of Cell 2's capacity unused. If the strategy instead performs remaining-capacity balancing (Strategy 2 in Section III-B), the two cells maintain equal remaining capacity after an initial balancing step, and the usable energy can be fully extracted, as shown in Fig. 3(b). However, since the Ah throughput of the two cells remains nearly identical, Cell 1 will still tend to reach EOL earlier over repeated cycling, wasting the remaining life of Cell 2. A similar conclusion applies to SOC balancing (Strategy 3 in Section III-B) in Fig. 3(c): although SOC balancing can improve energy utilization (enabled by the ability to bypass cells in the cell-level inverter topology), it does not preferentially reduce the stress on the weaker cell, so SOH disparity can grow over time and pack lifetime is not optimized.

In contrast, under SOH balancing (Fig. 3(d)), Cell 2 is deliberately used more than Cell 1, which reduces the SOH gap and can cause the two cells to reach EOL more simultaneously, minimizing wasted lifetime. The trade-off is that Cell 2 is depleted faster during discharge, which can reduce usable energy and hence driving range. Fig. 3(e) illustrates a more favorable strategy: it resembles SOH balancing during charging by allocating most Ah throughput to Cell 2, but it also performs remaining-capacity balancing during discharge so that no energy is left unused. This combined SOH + capacity balancing can therefore improve both long-term lifetime and energy utilization.

Finally, note that in Fig. 3(e) we apply SOH balancing during charging and remaining-capacity balancing during discharging, rather than the reverse. The primary reason is that battery degradation is SOC-dependent: high-SOC operation during charging can accelerate aging via both calendar aging mechanisms and charge-induced stress [27]. In contrast, although discharging contributes to cyclic aging, it also reduces SOC and can slow calendar aging, so its net impact on degradation is more uncertain [28]. Therefore, if SOC balancing is performed during charging while SOH balancing is attempted during discharge, the correct discharge allocation is ambiguous: it is unclear whether the weakest cells should be discharged more (to reduce time spent at high SOC) or less (to reduce cyclic-aging stress). Moreover, traction discharge power is exogenous and uncertain compared with scheduled charging, so implementing SOH balancing during discharge is considerably more complex. We therefore focus on SOH balancing during charging in this paper.

B. The proposed charging strategy

As discussed earlier, the core idea of "SOC–SOH-aware" control is to perform SOH balancing during charging and remaining-capacity balancing during discharge. The discharge strategy has been presented in Strategy 2. In this section, we focus on the charging controller, whose goal is to promote SOH balancing without creating an SOC distribution that cannot be rebalanced during the subsequent discharge.

1) Model simplifications: During charging, the controllable variables include the phase voltage, line current, and the cells' voltage-level assignments. These variables—especially the level assignments—can be updated at each control interval. Without simplification, the number of decision variables grows with the charging horizon and quickly becomes very large even for a small number of cells. Moreover, the level assignments are discrete (integer) decisions, which turn the problem into a large-scale mixed-integer optimization that is generally computationally intractable.

The primary purpose of the model simplifications is to reduce the degrees of freedom so that the number of decision variables does not grow with charging time. In Strategy 1, we showed that an assignment sequence can be generated once the desired final cell states are specified. Therefore, rather than optimizing the integer assignments at every time step, we optimize only the desired final states. For the phase voltage, we adopt a constant-voltage approximation and operate at the maximum permissible voltage magnitude (subject to cell and hardware voltage limits). As illustrated in Fig. 2, a higher phase-voltage magnitude corresponds to more cells connected in series, which allows a smaller line current for a fixed charging power. Since degradation often increases with C-rate, operating at the highest admissible phase voltage can reduce aging for a given charging-time requirement.

Finally, we assume a constant-current constant-voltage (CCCV) charging profile, which is widely used for EV batteries [29]. Specifically, cells are charged at a constant current (CC) until the maximum cell voltage reaches its limit; afterward, charging proceeds in a constant-voltage (CV) stage. Because the CV stage is largely determined by the cell states and voltage limit, the main continuous degree of freedom is the current magnitude in the CC stage.

With these simplifications, the remaining decision variables are:

- the target final added charge (capacity gain) of each cell, and
- the line current during the constant-current stage.

Although the number of decision variables is now small, the problem remains challenging because the charging-time constraint is difficult to express analytically. The total charging time depends on both the current profile and the required cell-wise capacity gains, and the current is not constant throughout the CCCV process. To further simplify the time constraint, we discretize the CCCV profile into multiple stages: Stage 0 represents constant-current charging, and Stages 1 through m represent constant-voltage charging. An example of this stage division is shown in Fig. 4.

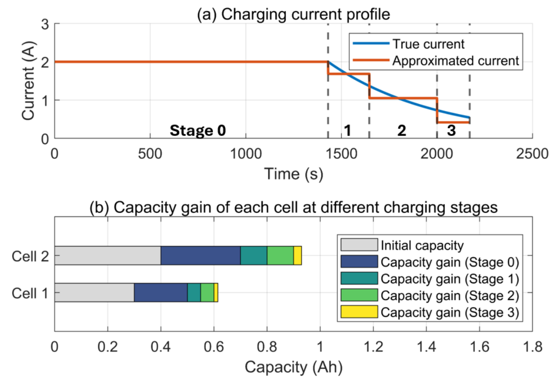

Fig. 4: The division of the charging stages.

In the example shown in Fig. 4, the CCCV charging is divided into four stages. The stage division is based on SOC, and the upper SOC boundary of Stage i is denoted by $SOC_{\max,i}$. We assume that the maximum admissible (normalized) C-rate is a function of SOC, denoted by $g(SOC)$. Stage 0 represents constant-current charging with a fixed line-current magnitude I_{cc}, which corresponds to a normalized C-rate I_{cc}/Q_{no}. In the current-envelope-limited setting, Stage 0 ends when this commanded C-rate reaches the admissible envelope, i.e.,

$$SOC_{\max,0} = g^{-1}\left(\frac{I_{cc}}{Q_{no}}\right). \tag{6}$$

The remaining SOC interval $[SOC_{\max,0}, 1]$ is then approximated by m constant-voltage sub-stages (Stages 1 through m).

Specifically, we choose the stage boundaries to uniformly partition $[SOC_{\max,0}, 1]$ in SOC:

$$SOC_{\max,i} = \begin{cases} g^{-1}\left(\dfrac{I_{cc}}{Q_{no}}\right), & i = 0, \\ SOC_{\max,0} + \dfrac{i}{m}\left(1 - SOC_{\max,0}\right), & i = 1, \ldots, m. \end{cases} \tag{7}$$

For each stage, we approximate the C-rate magnitude by a constant $I^{no}_{\text{avg},i}$ defined as

$$I^{no}_{\text{avg},i} = \begin{cases} \dfrac{I_{cc}}{Q_{no}}, & i = 0, \\ \dfrac{g(SOC_{\max,i-1}) + g(SOC_{\max,i})}{2}, & i = 1, \ldots, m, \end{cases} \tag{8}$$

where Q_{no} is the nominal cell capacity.

Within each stage, we assume the current magnitude (and thus the C-rate) is constant and equals $I^{no}_{\text{avg},i}$ (unit: C or

1/h). Under this approximation, the capacity gain delivered in each stage is proportional to that stage's duration. Let $Q_{p,i}$ denote the capacity gain (Ah) of cell p in Stage i. Then the total charging time can be expressed as a linear function of $\{Q_{p,i}\}$. By using $\{Q_{p,i}\}$ as decision variables, the time constraint becomes linear, substantially simplifying the optimization. This approximation becomes increasingly accurate as the number of stages m increases, at the cost of additional decision variables and computation time.

2) Formulation of the optimization problem: In summary, the model simplifications in the previous subsection are:

- The phase-voltage magnitude U_{phase} is set to the highest admissible constant (subject to voltage limits).
- The charging profile is approximated by CCCV charging, discretized into $m+1$ stages, with a constant line-current magnitude in each stage.
- The decision variables are the per-cell capacity gains in each stage, denoted by $Q_{i,j}$ for cell i and stage j. Note that the CC current magnitude I_{cc} is treated as a fixed parameter in the formulation below, and is optimized separately via a grid search.

We next define an objective function and constraints in terms of these decision variables. Ideally, the cost would reflect the capacity loss directly; however, this would require an explicit aging model, thereby reducing generality. Instead, we adopt a surrogate objective that discourages charging low-SOH cells and high-C-rate operation. For fixed U_{phase} and I_{cc}, we solve

$$\min_{\{Q_{i,j}\}} \sum_{i=1}^{n_1} \sum_{j=0}^{m} w_{i,j}\, Q_{i,j}, \tag{9}$$

where the weights are defined as

$$w_{i,j} = \begin{cases} \dfrac{1+\kappa I^{no}_{\text{avg},j}}{(SOH_i - SOH_{\text{cell,EOL}})^p}, & SOH_i > SOH_{\text{cell,EOL}} + \varepsilon, \\ \left(1+\kappa I^{no}_{\text{avg},j}\right) M, & SOH_i \le SOH_{\text{cell,EOL}} + \varepsilon. \end{cases} \tag{10}$$

Here, SOH_i denotes the available SOH estimate of cell i, and $SOH_{\text{cell,EOL}}$ is the cell-level EOL threshold. The exponent p determines how sharply the penalty increases as a cell approaches EOL. In the default case, we select $p = 2$, $\kappa = 0.1$, $\varepsilon = 10^{-3}$, and $M = \max\{10^6, 10\varepsilon^{-p}\}$. The reference choice $p = 2$ yields a moderate degree of nonlinear prioritization over healthier cells. Alternative values of p and κ are evaluated in Section V-C. Once a cell reaches $SOH_{\text{cell,EOL}}$, the large M strongly discourages further charging, and the cell is bypassed.

The constraints of the optimization problem are listed below:

1) **Nonnegativity of stage-wise allocations.**

$$Q_{i,j} \ge 0, \qquad \forall i, \forall j. \tag{11}$$

2) **Fixed total added charge content.**

$$\sum_{i=1}^{n_1} \sum_{j=0}^{m} Q_{i,j} = Q_{\text{final,sum}} - Q_{\text{initial,sum}}, \tag{12}$$

where $Q_{\text{final,sum}}$ and $Q_{\text{initial,sum}}$ denote the target and initial *total charge contents* (Ah) summed across all n_1 cells, respectively.

3) **Charging-time limit.**

$$\sum_{j=0}^{m} \frac{\sum_{i=1}^{n_1} Q_{i,j}}{Q_{no}\, I^{no}_{\text{avg},j} \sum_{k=1}^{n_1} d_{k,j}} \le t_{\text{total}}, \tag{13}$$

where t_{total} is the charging-time limit (hours), and $d_{i,j}$ is the duty cycle of voltage level i in stage j. For sinusoidal modulation, we approximate

$$d_{i,j} = \frac{2}{\pi} \arccos\left(\min\left(\frac{(2i-1)U_{\text{ter},j}}{2U_{\text{phase}}}, 1 \right) \right), \tag{14}$$

where $U_{\text{ter},j}$ is the representative per-cell terminal voltage in stage j and U_{phase} is the (fixed) phase-voltage magnitude. For DC modulation, we use

$$d_{i,j} = \begin{cases} 0, & (i-1)U_{\text{ter},j} > U_{\text{phase}}, \\ 1, & iU_{\text{ter},j} < U_{\text{phase}}, \\ \dfrac{U_{\text{phase}} - (i-1)U_{\text{ter},j}}{U_{\text{ter},j}}, & \text{otherwise}. \end{cases} \tag{15}$$

Strictly speaking, terminal voltages vary across cells and time. We approximate the stage-wise terminal voltage by the OCV at an average SOC plus an average ohmic drop:

$$U_{\text{ter},j} \approx f_{OCV}\left(\frac{Q_{\text{final,sum}} + Q_{\text{initial,sum}}}{2\sum_{i=1}^{n_1} Q_{\max,i}} \right) + Q_{no}\, I^{no}_{\text{avg},j} R_0, \tag{16}$$

where $f_{OCV}(\cdot)$ is the OCV–SOC map, $Q_{\max,i}$ is the current maximum capacity of cell i, and R_0 is a representative source-side resistance. A common value is used to ensure that the duty-cycle matrix $d_{i,j}$ is fixed and independent of the decision variables. In practice, R_0 can be selected as an upper bound over the active cells, making the ohmic components of the charging and discharge voltage estimates conservative.

4) **Per-stage SOC upper bounds (voltage-limit proxy).** At the end of each stage k, the SOC of each cell must not exceed $SOC_{\max,k}$ defined in (7):

$$Q_{\text{initial},i} + \sum_{j=0}^{k} Q_{i,j} \le Q_{\max,i}\, SOC_{\max,k}, \quad \forall i, \forall k. \tag{17}$$

5) **Stage-wise achievability under LSPWM.** According to Theorem 1, for each stage j, the target allocation vector $\{Q_{i,j}\}_{i=1}^{n_1}$ must satisfy

$$\frac{\sum_{i=1}^{k} Q_{(i),j}}{\sum_{i=1}^{n_1} Q_{i,j}} \le \frac{\sum_{\ell=1}^{k} d_{\ell,j}}{\sum_{\ell=1}^{n_1} d_{\ell,j}}, \quad \forall k, \forall j, \tag{18}$$

where $Q_{(1),j} \ge \cdots \ge Q_{(n_1),j}$ denotes the sorted values of $\{Q_{i,j}\}_{i=1}^{n_1}$. Because level assignments can be permuted independently across stages and each stage is treated as a constant-duty allocation subproblem, imposing (18) for each stage is sufficient under the stage-wise approximation.

A direct implementation of (18) involves sorting. We instead enforce it using an epigraph formulation. Introduce auxiliary variables $t_{k,j} \in \mathbb{R}$ and $s_{i,k,j} \ge 0$ such that

$$s_{i,k,j} \ge Q_{i,j} - t_{k,j}, \qquad s_{i,k,j} \ge 0, \qquad \forall i, \forall k, \forall j. \tag{19}$$

For any fixed (k, j), minimizing $k\, t_{k,j} + \sum_i s_{i,k,j}$ over $t_{k,j}$ yields $\sum_{i=1}^{k} Q_{(i),j}$. Therefore, (18) is equivalently enforced by requiring the *existence* of $(t_{k,j}, s_{i,k,j})$ satisfying (19) and

$$k\, t_{k,j} + \sum_{i=1}^{n_1} s_{i,k,j} \le \left(\frac{\sum_{\ell=1}^{k} d_{\ell,j}}{\sum_{\ell=1}^{n_1} d_{\ell,j}} \right) \sum_{i=1}^{n_1} Q_{i,j}, \qquad \forall k, \forall j. \tag{20}$$

6) **SOH-consistent ordering of the final stored charge (structural constraint).** Cells with higher SOH are required to have higher stored charge by the end of charging:

$$Q_{\text{initial},i} + \sum_{k=0}^{m} Q_{i,k} \leq Q_{\text{initial},j} + \sum_{k=0}^{m} Q_{j,k}, \text{if } SOH_i < SOH_j. \tag{21}$$

This constraint aligns the solution with the intended SOH-aware usage pattern and facilitates the discharge-feasibility constraint below.

7) **Sufficient condition for full utilization during subsequent discharge.** We consider the target discharge end state $Q_{\text{end},i} = 0$ for all cells; hence the required discharge amounts are $\Delta Q_i = Q_{\text{final},i}$, where $Q_{\text{final},i} \triangleq Q_{\text{initial},i} + \sum_{k=0}^{m} Q_{i,k}$. Applying Theorem 1 yields the sufficient condition

$$\frac{\sum_{i=1}^{k} Q_{\text{final},(i)}}{\sum_{i=1}^{n_1} Q_{\text{final},i}} \leq \frac{\sum_{i=1}^{k} d_i'}{\sum_{i=1}^{n_1} d_i'}, \qquad \forall k \in \{1, \ldots, n_1\}, \tag{22}$$

where $Q_{\text{final},(1)} \geq \cdots \geq Q_{\text{final},(n_1)}$ are the sorted final charge contents and $\{d_i'\}$ is a representative discharge duty-cycle pattern computed from average discharge voltage and phase-voltage magnitude. Similar to (14), we approximate

$$d_i' = \frac{2}{\pi} \arccos\left(\min\left(\frac{(2i-1)U_{\text{dis,ter}}}{2U_{\text{dis,phase}}}, 1 \right) \right), \tag{23}$$

with

$$U_{\text{dis,ter}} \approx f_{OCV}\left(\frac{Q_{\text{final,sum}}}{2\sum_{i=1}^{n_1} Q_{\max,i}} \right) - Q_{no}\, I_{\text{dis,avg}}^{no}\, R_0, \tag{24}$$

where $I_{\text{dis,avg}}^{no}$ is the average C rate during discharge. For a more conservative design, it may be replaced by the maximum allowable discharge C-rate. Because (21) fixes the ordering of the final stored charge, (22) can be imposed directly as a linear constraint.

For any fixed U_{phase} and I_{cc}, the problem (9) subject to Constraints 1–7 is a linear program (LP) in the decision variables $\{Q_{i,j}\}$ and can be solved efficiently. The overall procedure for selecting U_{phase}, I_{cc}, and $\{Q_{i,j}\}$ is summarized in Fig. 5. We initialize U_{phase} and I_{cc} at their maximum admissible values $(U_{\text{phase,max}}, I_{cc,\max})$. A larger U_{phase} increases the minimum number of active voltage levels and therefore reduces the achievable imbalance; accordingly, we decrease U_{phase} until the LP becomes feasible. After U_{phase} is determined, we enumerate candidate values of I_{cc} and select the pair $(I_{cc}, \{Q_{i,j}\})$ that minimizes the objective in (9).

In some cases, the driver may want to charge the vehicle to a specific SOC level as quickly as possible. This can be achieved by skipping the step of finding the optimal I_{cc} and setting I_{cc} to the cell's maximum charging current. In other words, only the first two steps in Figure 5 are required, and the rest can be skipped.

After the optimization problem is solved, the phase voltage and line current are set as the optimized values. In implementation, the level assignments are adjusted online via Strategy 1, which provides robustness to deviations from the average-voltage approximations used in the LP.

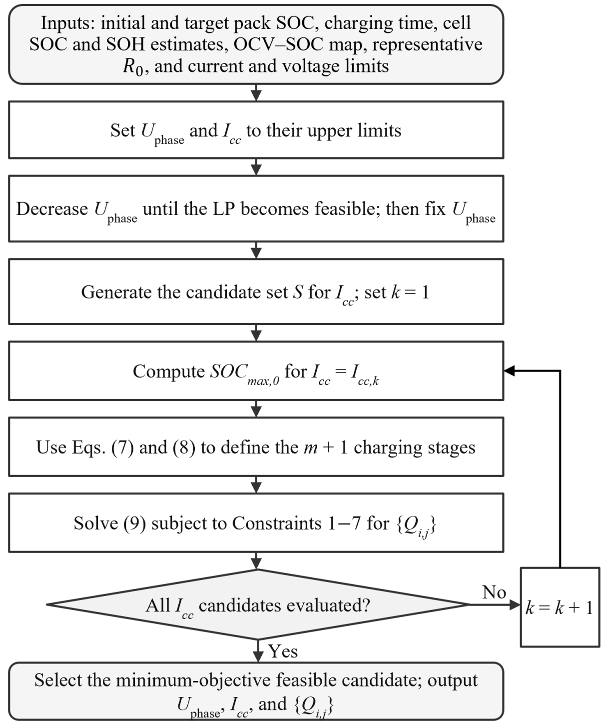

Fig. 5: The flow chart of the optimization algorithm.

3) Computational complexity and reduced formulation: In the LP formulation, Constraint 5 is the main source of computational growth. Its epigraph representation introduces $n_1(m+1)$ variables $t_{k,j}$ and $n_1^2(m+1)$ variables $s_{i,k,j}$. The full LP therefore contains $(m+1)(n_1^2+2n_1)$ decision variables and $(m+1)(n_1^2+2n_1)+2n_1-1$ inequality constraints, so its size grows quadratically with n_1.

The charging optimization is performed once at the beginning of each AC charging session, using the most recent per-cell SOC and SOH estimates. The controller therefore requires state estimates to be refreshed at most once per charging session. For $n_1 = 20$ and $m = 6$, each full LP contains at most 3080 variables, 3119 inequalities, and one equality. On a PC with an Intel Core Ultra 9 275HX processor at 2.70 GHz, the median LP solve time was 60.3 ms. A median of 55 LPs was solved per session, giving median and maximum session-level optimization times of 3.95 s and 10.13 s, respectively. The stored LP problem arrays occupied at most 0.415 MB.

For larger control groups, we introduce a reduced formulation. Let $\pi_1, \ldots, \pi_{n_1}$ order the cells from highest to lowest SOH. The reduced formulation requires $Q_{\pi_1,j} \geq \cdots \geq Q_{\pi_{n_1},j}$ in every stage j. The k largest stage allocations in Constraint 5 are then known in advance, so the constraint can be imposed directly without $t_{k,j}$ and $s_{i,k,j}$. The reduced LP contains $n_1(m+1)$ variables and $(3n_1-2)(m+1)+2n_1-1$ inequalities. This stage-wise ordering restriction makes the feasible set more conservative, while retaining the same objective and all other constraints.

For $n_1 = 20$ and $m = 6$, the reduced LP contains 140 variables, 445 inequalities, and one equality. With the same voltage and current search, the median LP solve time was 6.90 ms, and the median and maximum session-level optimization times were 0.473 s and 0.895 s, respectively. The stored LP problem arrays occupied at most 0.111 MB.

V. Simulation study

A. Simulation setup

The proposed controller is evaluated through simulation rather than pack-level experiments for two reasons. First, the objective of this paper is to examine a new cell-level energy-throughput control concept whose benefit depends on long-term interactions among cell aging heterogeneity, EOL definition, charging behavior, and discharge utilization. A pack-level aging experiment would require many cells, long test durations, and repeated trials to distinguish the effect of the controller from cell-to-cell variability. Second, a single experimental setup would validate only one combination of chemistry, usage profile, heterogeneity level, and EOL criterion, whereas the central question here is whether the proposed control principle remains beneficial across a range of plausible operating conditions. Therefore, we use simulation as a controlled evaluation platform that enables systematic comparison across aging models and parameter settings. Experimental validation on hardware is an important next step, but is beyond the scope of this first concept-level study.

Specifically, we perform two types of simulations. First, a drive-cycle simulation examines how each cell's SOC evolves during charging and discharging. The discharge current profile is extracted from a real-world driving test and repeatedly used until the pack energy is depleted. The results (similar to Fig. 3(e)) illustrate how the proposed strategy intentionally creates non-uniform utilization during charging (i.e., different cells receive different effective charge throughput) while still enabling full energy utilization during discharge.

Second, a long-term aging simulation quantifies the lifetime improvement of the proposed strategy relative to conventional SOC balancing. To construct realistic usage conditions, we use charging records collected from a 2021 Tesla Model 3 over approximately 2.5 years, comprising 255 cycles, as shown in Fig. 6(a). In the simulation, these profiles are repeated until the battery pack reaches EOL. Among these charging events, 82 are DC fast charging with an average C rate greater than 0.2 C, and the remaining 173 are AC charging; Fig. 6(b) shows the corresponding distributions of the session-average C-rate and the charging duration for the two charging types. Since the purpose of DC fast charging is typically to achieve sufficient capacity as quickly as possible, we always use the highest available voltage and current to maximize charging speed and suppress SOH balancing. Meanwhile, for each AC charging event, the estimated initial and target pack SOCs and the charging duration (capped at 10 h) are provided as inputs to the optimization algorithm in Fig. 5, whose output determines the cell-level charging commands. During discharge, we apply remaining-capacity balancing, and each discharge terminates when the pack's SOC reaches the recorded initial SOC of the subsequent charging session. This procedure is repeated until the pack reaches the EOL criterion defined in Section II-B.

Note that, in both simulations, the proposed strategy balances SOH only within each phase. Therefore, for brevity, we report results only for the first phase of the three.

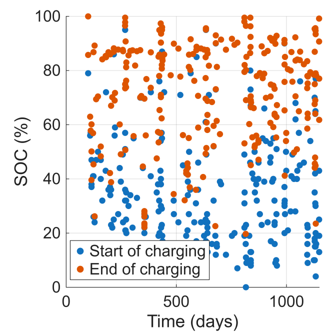

(a) The charging records used for the long-term aging simulation.

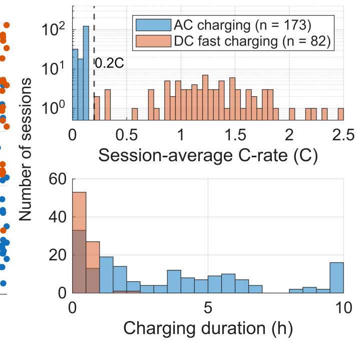

(b) Distributions of the session-average C-rate (top) and the charging duration (bottom) across the recorded charging events.

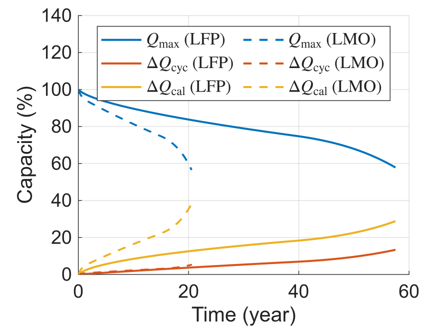

(c) Typical aging curves of two types of lithium-ion batteries.

Fig. 6: Long-term aging simulation setups.

The default specifications of the battery pack and chargers used in the simulations are as follows. Note that these specifications are not fixed across all cases; the impact of varying them is discussed in Section V-C. By default, each cell has a nominal capacity of 2.3 Ah, and each phase comprises 20 cells. The representative resistance used in the controller is $R_0 = 0.01\,\Omega$. The same value is used in the reference-case diagnostics. During relaxation, the surface temperature of each cell is assumed to be 25°C. During driving and charging, the surface temperature is modeled as a cell-to-cell random variation in the default case. Namely, for each cell, T is sampled once from a Gaussian distribution with mean 35°C and standard deviation 2°C, and the sampled value is held constant over time. In Section V-C, we also include a first-order thermal-feedback case where each cell's temperature rise during a charge or discharge event is proportional to its Ah throughput, while the cell-average rise remains 10 °C. During charging, the maximum allowable C-rate $I^{no}_{\max}$ is modeled as a function of SOC. By fitting the charging data of a Tesla Model 3 available on EVKX.net, we obtain the following empirical function:

$$I^{no}_{\max} = 2.6963 - 2.5795 \cdot SOC \tag{25}$$

Two lithium-ion chemistries are considered: LFP and LMO. The OCV–SOC relationships are adopted from [30], [31], and the baseline aging models are adopted from [32], [33]. In these models, calendar aging depends on SOC, temperature, and storage time, whereas cycling aging depends on temperature,

This article has been accepted for publication in IEEE Transactions on Transportation Electrification. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TTE.2026.3725087

C-rate, Ah throughput, depth of discharge, and average SOC. The detailed OCV functions, aging equations, and parameter values are provided in the supplementary material available in the online code repository.

The baseline aging models are used only for evaluating the proposed control strategy after optimization; they are not embedded in the controller. Therefore, the proposed strategy remains degradation-model-agnostic with respect to battery degradation. However, the baseline models have two limitations in the present pack-level setting. First, they do not explicitly account for cell-to-cell aging heterogeneity within a pack. Second, they do not explicitly represent the accelerated late-life degradation observed after the onset of a capacity-fade knee [34]. We therefore augment the baseline models with a cell-specific aging-rate factor and a tunable late-life acceleration term. Specifically, for either chemistry, the short-term degradation of cell i is computed as

$$\Delta SOH_i = \gamma_i \Big[1 + a_{\mathrm{knee}} \max\{SOH_{\mathrm{knee}} - SOH_i,\, 0\} \Big] \Delta SOH_{\mathrm{base}}. \tag{26}$$

Here, $\Delta SOH_{\mathrm{base}}$ is the degradation predicted by the corresponding chemistry-specific aging model. The factor $\gamma_i \sim \mathcal{N}(1, \sigma_\gamma^2)$ is sampled once for each cell and held constant, representing persistent cell-to-cell variation in aging rate. Unless otherwise stated, $\sigma_\gamma = 0.10$, $SOH_{\mathrm{knee}} = 75\,\%$, and $a_{\mathrm{knee}} = 25$, with SOH represented as a fraction in Eq. (26). The resulting late-life multiplier is 1 at $75\,\%$ SOH, 2.25 at $70\,\%$ SOH, and 4.75 at $60\,\%$ SOH. We use $a_{\mathrm{knee}} = 25$ as a moderate reference value and evaluate no added acceleration, the stronger value $a_{\mathrm{knee}} = 50$, and an earlier knee onset in Section V-C.

B. Simulation results

The drive-cycle results are shown in Fig. 7. For visualization, each phase contains 10 LFP cells, ranked by SOH (cell i has the ith highest SOH). Panels (a) and (b) use the proposed strategy. In both cases, the pack starts at $0\,\%$ SOC, is charged to $70\,\%$, and is then fully discharged. Under the charging-time and pack-SOC constraints, the proposed strategy allocates different charging amounts to different cells. The key difference between the two cases is the available charging time: Fig. 7(a) uses a shorter charging window than Fig. 7(b). As shown, cells with higher SOH are charged more than those with lower SOH, which promotes long-term SOH balancing. Moreover, the remaining capacity of all cells is balanced within the first 25% of the discharge, indicating that the strategy does not reduce the usable energy (and thus range) in this example. Panel (c) repeats the 1.5-h case but removes Constraint 7, thereby representing SOH balancing without a subsequent-discharge guarantee.

Figure 7(c) illustrates the range penalty of unconstrained SOH balancing. This case retains the SOH-weighted objective and Constraints 1–6 but omits the condition that preserves full utilization during subsequent discharge. It also applies SOH-oriented level assignment during the subsequent drive cycle. Cell 10 receives no charge and remains at $0\,\%$ SOC. During the early part of the discharge, the requested line voltage can be supplied using the other cells, so Cell 10 remains bypassed.

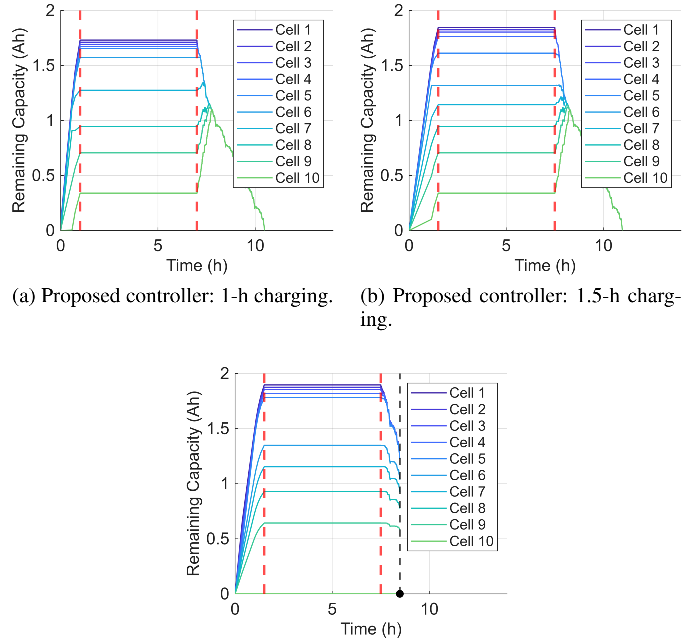

(a) Proposed controller: 1-h charging.

(b) Proposed controller: 1.5-h charging.

(c) SOH balancing.

Fig. 7: Drive-cycle comparison of the proposed controller and SOH balancing. Panels (a) and (b) use the proposed controller with charging windows of 1 and 1.5 h, respectively. Panel (c) uses the same 1.5-h charging window, removes the discharge-feasibility constraint, and applies SOH balancing during the subsequent drive cycle. The red dashed lines mark the end of charging and the beginning of discharge. In panel (c), the requested line voltage can initially be supplied without using Cell 10, which remains at $0\,\%$ SOC. At $t = 8.480$ h, the other cells can no longer satisfy the voltage requirement alone. Continuing the requested discharge would therefore require Cell 10 to fall below $0\,\%$ SOC, although the pack SOC remains $49.75\,\%$.

At $t = 8.480$ h, the other cells can no longer satisfy the voltage requirement alone. Continuing the requested discharge would therefore require Cell 10 to fall below $0\,\%$ SOC, even though the pack SOC remains $49.75\,\%$. Thus, SOH balancing alone can leave substantial energy inaccessible and sacrifice driving range. The proposed controller avoids this trade-off by combining the discharge-feasibility constraint during charging with remaining-capacity balancing during the subsequent drive cycle.

Comparing Fig. 7(a) and (b), we observe that the SOC variance across cells at the end of charging is larger when the charging time is longer. This occurs because the CV stage is relatively slow: with a longer charging window, the controller can spend more time in the CV regime and push healthier cells to higher SOC while bypassing less healthy cells. Since the target pack SOC is fixed, this increases the cell-to-cell SOC dispersion and, in turn, enables more aggressive SOH balancing when more charging time is available.

Another observation from Fig. 7(a) and (b) is that several cells share identical SOC trajectories during charging. This is a consequence of the CV voltage limit: the pack charging current is constrained by the cell that reaches the voltage limit first (typically the cell with the highest SOC), which encourages the controller to equalize SOC among a subset of cells to increase the allowable current and accelerate charging. In contrast, during discharge, the strategy balances remaining

This article has been accepted for publication in IEEE Transactions on Transportation Electrification. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TTE.2026.3725087

capacity, so cells tend to maintain similar remaining capacities rather than identical SOC values.

In the long-term aging simulations, the baseline controller is the SOC-balancing strategy (Strategy 3 in Section III-B), which equalizes cell SOC by assigning the highest-duty voltage levels to lower-SOC cells during charging and to higher-SOC cells during discharge. SOH balancing is not considered as a long-term aging baseline, because it can render the requested discharge infeasible before the stored pack energy is exhausted, as exemplified in Fig. 7(c). To provide a stronger range-oriented comparison, we also consider ideal remaining-capacity balancing. This benchmark assumes exact per-cell capacity information, makes the final stored charge contents as equal as possible subject to the cell SOC limits, and applies remaining-capacity balancing during discharge. Additionally, this ideal benchmark is not restricted by the LSPWM achievability condition and therefore represents an optimistic range-oriented comparison. For a fair comparison, all three long-term strategies are evaluated using identical charging records, cell realizations, and aging assumptions. Since the proposed strategy relies on per-cell SOH information and practical SOH estimators are imperfect, the reference proposed controller is evaluated with independent, zero-mean Gaussian SOH-estimation errors having a standard deviation of 2 percentage points. Figure 8 compares SOC balancing, remaining-capacity balancing, and the reference proposed controller for both chemistries.

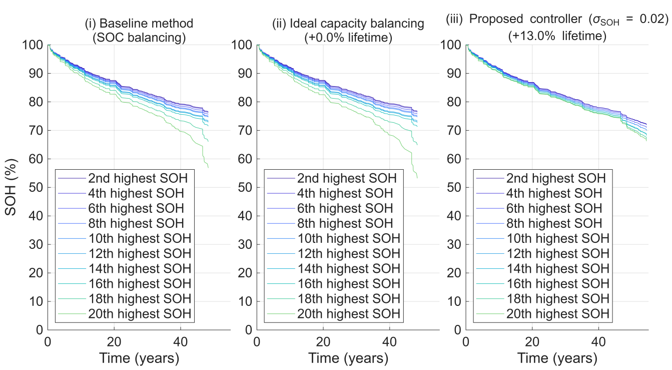

(a) LFP battery

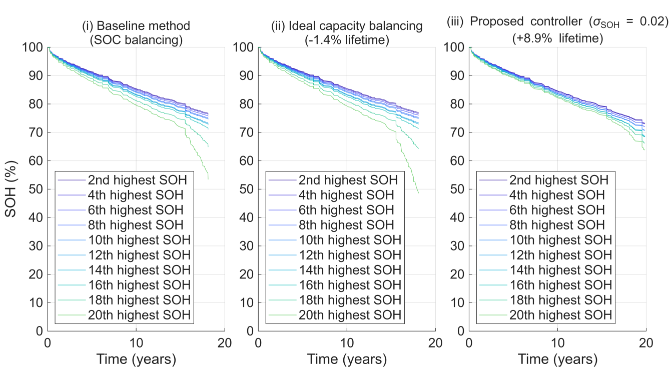

(b) LMO battery

Fig. 8: Long-term SOH trajectories under (i) SOC balancing, (ii) ideal remaining-capacity balancing with exact per-cell capacity information, and (iii) the proposed controller with zero-mean SOH-estimation noise having a standard deviation of 0.02, for (a) LFP and (b) LMO batteries.

Because pack EOL is evaluated at the end of each recorded charging session, ideal remaining-capacity balancing and SOC balancing reach EOL in the same session for LFP. For LMO, ideal remaining-capacity balancing yields a 1.4 % shorter lifetime than SOC balancing. Meanwhile, the proposed controller extends lifetime by 13.0 % for LFP and 8.9 % for LMO. This result confirms that remaining-capacity balancing alone does not explain the lifetime extension. The additional benefit of the proposed controller comes from deliberately shifting charging throughput away from lower-SOH cells while preserving subsequent usable energy.

We also implement the reduced formulation proposed in Section IV-B3. The reduced formulation yields lifetime improvements of 13.0 % for LFP and 8.7 % for LMO, compared with 13.0 % and 8.9 % under the full formulation. Thus, the stage-wise ordering restriction has little effect on the predicted lifetime benefit in the reference case.

C. Sensitivity analysis

Because the controller is degradation-model-agnostic, its effectiveness should not depend on one specific aging law, one objective-function design, one EOL convention, or perfect state and parameter information. We therefore vary the late-life aging assumptions, the SOH-weight exponent and C-rate coefficient, the cell- and pack-EOL thresholds, the SOH and SOC estimates, the OCV and resistance models, the temperature model, and selected usage conditions. Each row of Table I changes only the listed setting from the reference case.

Table I shows that the lifetime benefit is insensitive to most controller-design and parameter-mismatch choices. Varying p from 1 to 3 or κ from 0.05 to 0.20 produces almost no change. The OCV bias and the increase in representative resistance from 0.01 to $0.05\,\Omega$ also leave the reported benefit unchanged over the tested range. Larger SOH noise, SOC bias, throughput-coupled heating, and a 50 % DC-fast-charging share reduce the benefit only moderately, and it remains positive for both chemistries. The fast-charging result is expected because the present strategy suppresses SOH balancing during those sessions. Limited SOH-aware allocation during the CV taper is a possible extension.

The largest variations arise from the aging and EOL assumptions. The key distinction is whether cell retirement becomes limiting before pack EOL. In the coupled $80\,\%/80\,\%$ case, a cell is removed as soon as its SOH reaches $80\,\%$. With $n_1 = 20$, retiring one cell near this threshold reduces the modeled pack SOH by approximately $0.8/20 = 0.04$. Thus, the first weak-cell retirement can bring the pack close to, or below, its $80\,\%$ EOL threshold. The proposed controller delays this event by maintaining a more uniform SOH distribution, whereas SOC balancing allows the weakest cells to reach EOL earlier. This amplifies the lifetime gain to $36.7\,\%$ for LFP and $17.3\,\%$ for LMO. By contrast, in the $80\,\%/60\,\%$ case, pack EOL occurs before cell retirement becomes relevant, so balancing the SOH distribution provides no benefit and may slightly reduce lifetime. The near-zero benefit without late-life acceleration reflects the same principle: the controller is most

TABLE I: Lifetime change of the proposed controller relative to SOC balancing under different assumptions. Each row changes only the listed setting from the reference case.

Scenario	**LFP (%)**	**LMO (%)**
Reference case	13.0	8.9
Aging-model assumptions		
Calendar aging $\times 2$	8.8	9.4
Cyclic aging $\times 2$	9.4	9.4
No late-life acceleration ($a_{\text{knee}} = 0$)	0.0	1.6
Stronger late-life acceleration ($a_{\text{knee}} = 50$)	9.0	6.8
Earlier knee onset ($SOH_{\text{knee}} = 85\,\%$)	5.3	9.1
Throughput-coupled temperature rise: $\Delta T_i \propto \lvert \Delta Ah_i \rvert$	11.5	8.7
Objective-function design		
SOH-weight exponent $p = 1$	13.0	8.7
SOH-weight exponent $p = 3$	13.0	8.9
C-rate weight $\kappa = 0.05$	13.0	8.9
C-rate weight $\kappa = 0.20$	13.0	8.9
Cell and pack EOL definitions		
Lower cell-EOL threshold (pack/cell: $70\,\%/50\,\%$)	8.3	6.2
Coupled EOL thresholds (pack/cell: $70\,\%/70\,\%$)	6.5	8.3
Higher pack-EOL threshold (pack/cell: $80\,\%/60\,\%$)	0.0	-2.3
Coupled high-EOL thresholds (pack/cell: $80\,\%/80\,\%$)	36.7	17.3
Estimation and parameter mismatch		
Exact SOH information	13.0	9.7
SOH noise: $\sigma_{\text{SOH}} = 0.05$	11.5	8.7
Persistent SOH bias: $+0.05$	13.0	8.9
Persistent SOC bias: $+0.05$	11.5	8.7
OCV-model bias: -20 mV	13.0	8.9
Conservative representative resistance: $R_{0,\text{rep}} = 0.05\,\Omega$	13.0	8.9
Different use cases		
Higher aging-rate spread: $\sigma_\gamma = 0.15$	12.7	4.5
DC fast-charging share: $50\,\%$	10.4	7.8

Lifetime is resolved at the granularity of complete charging sessions; distinct perturbations can therefore yield identical percentage changes. *Reference case:* $n_1 = 20$, $SOH_{\text{pack,EOL}} = 70\,\%$, $SOH_{\text{cell,EOL}} = 60\,\%$, $SOH_{\text{knee}} = 75\,\%$, $a_{\text{knee}} = 25$, $\sigma_\gamma = 0.10$, $p = 2$, and $\kappa = 0.1$. The SOH estimates used in Eq. (10) contain additive, zero-mean Gaussian errors with a standard deviation of 0.02 on the 0–1 SOH scale. Errors are redrawn independently for each cell at every charging session. The SOH-bias case retains this random error, with the bias held fixed throughout the simulation. The SOC- and OCV-bias cases affect controller inputs only. Physical aging and EOL are evaluated using the true states. In the conservative-resistance case, $R_0 = 0.05\,\Omega$ is used in both the controller and the simulation. In the throughput-coupled temperature case, each cell's operating-temperature rise is proportional to its Ah throughput during the corresponding charge or discharge event, while the cell-average rise is fixed at $10\,^\circ$C.

valuable when weak-cell retirement, rather than average aging alone, limits usable pack life.

VI. Conclusions

In this paper, we propose a SOC–SOH-aware control strategy to extend the range and service life of EV battery packs equipped with a cell-level inverter powertrain. We focus on applying the strategy during charging, particularly during AC charging events. The strategy does not require an explicit aging model or a forecast of the next trip. Practical implementation nevertheless requires cell-level switching and standard BMS information on cell states and electrical limits. In the reference case with 2-percentage-point SOH-estimation noise, the proposed strategy extends pack lifetime by 13.0 % for LFP and 8.9 % for LMO relative to SOC balancing, while maintaining comparable usable energy for subsequent driving. The sensitivity analysis quantifies the effects of late-life aging, objective-function design, EOL definitions, and state and parameter uncertainty. It also shows that the proposed health-prioritized allocation provides little or no benefit when pack EOL occurs before any cell reaches the cell-level EOL threshold. The proposed health-prioritization principle may also be adapted to conventional active-balancing hardware, although the lower power of auxiliary balancing converters would make SOH equalization slower [35], [36].

Although the present study does not include pack-level hardware aging experiments, the simulations provide a broad evaluation of the control concept across multiple aging models and operating assumptions. The analysis assumes frequent level reassignment and uses stage-wise voltage and simplified thermal models. Converter losses and other hardware-level effects are not represented. Future work will incorporate dynamic electrothermal and converter-loss models and validate the strategy on representative hardware.

References

[1] K. Kandasamy, M. Vilathgamuwa, and K. J. Tseng, "Inter-module state-of-charge balancing and fault-tolerant operation of cascaded H-bridge converter using multi-dimensional modulation for electric vehicle application," *IET Power Electronics*, vol. 8, no. 10, pp. 1912–1919, 2015.

[2] T. Wu, F. Ji, L. Liao, and C. Chang, "Voltage-SOC balancing control scheme for series-connected lithium-ion battery packs," *Journal of Energy Storage*, vol. 25, p. 100895, 2019.

[3] F. Gemma, G. Tresca, A. Formentini, and P. Zanchetta, "Balanced charging algorithm for CHB in an EV powertrain," *Energies*, vol. 16, no. 14, p. 5565, 2023.

[4] G. Tresca, A. Formentini, J. Riccio, N. Anglani, and P. Zanchetta, "A reconfigurable cascaded multilevel converter for EV powertrain," *IEEE Transactions on Industry Applications*, vol. 60, no. 2, pp. 3332–3344, 2024.

[5] G. Tresca and P. Zanchetta, "AC direct charging for electric vehicles via a reconfigurable cascaded multilevel converter," *Energies*, vol. 17, no. 10, p. 2428, 2024.

[6] Stellantis, "World premiere: First prototype vehicle featuring breakthrough intelligent battery integrated system begins real-world testing," https://www.media.stellantis.com/em-en/corporate/press/world-premiere-first-prototype-vehicle-featuring-breakthrough-intelligent-battery-integrated-system-begins-real-world-testing, Sep. 2025, accessed: Jul. 12, 2026.

[7] F. Roy, A. Dittrick, B. Revol, D. Labrousse, and T. Peuchant, "Method for controlling a battery allowing control of a homogenous current to the cells of a power line," U.S. Patent 12 212 166, Jan. 28, 2025.

[8] T. Baumhöfer, M. Brühl, S. Rothgang, and D. U. Sauer, "Production caused variation in capacity aging trend and correlation to initial cell performance," *Journal of Power Sources*, vol. 247, pp. 332–338, 2014.

[9] J. Lin, X. Liu, S. Li, C. Zhang, and S. Yang, "A review on recent progress, challenges and perspective of battery thermal management system," *International Journal of Heat and Mass Transfer*, vol. 167, p. 120834, 2021.

[10] Z. Xia and J. A. Abu Qahouq, "State-of-charge balancing of lithium-ion batteries with state-of-health awareness capability," *IEEE Transactions on Industry Applications*, vol. 57, no. 1, pp. 673–684, 2021.

[11] G. Kim, S. Lee, S. Chun, and Y. Kwon, "Analysis of battery lifespan improvement with SOC battery balancing technique," *Journal of Energy Storage*, vol. 137, p. 118543, 2025.

[12] A. Pröbstl, S. Park, S. Narayanaswamy, S. Steinhorst, and S. Chakraborty, "SOH-aware active cell balancing strategy for high power battery packs," in *2018 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2018, pp. 431–436.

[13] Y. Yu, B. Wang, T. Lv, and X. Chen, "Rapid energy management strategy for multi-lithium battery systems considering synergistic balancing of SOC and SOH under time-delay environments," *Energy*, vol. 336, p. 138429, 2025.

[14] S. Jiang and Z. Song, "A review on the state of health estimation methods of lead-acid batteries," *Journal of Power Sources*, vol. 517, p. 230710, 2022.

[15] M. Etxandi-Santolaya, L. C. Casals, T. Montes, and C. Corchero, "Are electric vehicle batteries being underused? a review of current practices and sources of circularity," *Journal of Environmental Management*, vol. 338, p. 117814, 2023.

[16] E. Martinez-Laserna, I. Gandiaga, E. Sarasketa-Zabala, J. Badeda, D.-I. Stroe, M. Swierczynski, and A. Goikoetxea, "Battery second life: Hype, hope or reality? a critical review of the state of the art," *Renewable and Sustainable Energy Reviews*, vol. 93, pp. 701–718, 2018.

[17] S. Saxena, C. Le Floch, J. MacDonald, and S. Moura, "Quantifying EV battery end-of-life through analysis of travel needs with vehicle powertrain models," *Journal of Power Sources*, vol. 282, pp. 265–276, 2015.

[18] "USABC Electric Vehicle Battery Test Procedures Manual, Revision 2," United States Advanced Battery Consortium, Tech. Rep. DOE/ID-10479-Rev.2, Jan. 1996.

[19] J. P. Christophersen, *Battery Test Manual for Electric Vehicles, Revision 3*, Idaho National Laboratory, Jun. 2015, Tech. Rep. INL/EXT-15-34184. [Online]. Available: https://www.osti.gov/biblio/1186745

[20] California Air Resources Board, *Section 1962.8 - Warranty Requirements for Zero-Emission and Batteries in Plug-in Hybrid Electric 2026 and Subsequent Model Year Passenger Cars and Light-Duty Trucks*, California Code of Regulations, Mar. 2024.

[21] ——, *Section 1962.5 - Data Standardization Requirements for 2026 and Subsequent Model Year Light-Duty Zero Emission Vehicles and Plug-in Hybrid Electric Vehicles*, California Code of Regulations, Mar. 2024.

[22] L. Zhong, C. Zhang, Y. He, and Z. Chen, "A method for the estimation of the battery pack state of charge based on in-pack cells uniformity analysis," *Applied Energy*, vol. 113, pp. 558–564, 2014.

[23] M. B. Latran and A. Teke, "Investigation of multilevel multifunctional grid connected inverter topologies and control strategies used in photovoltaic systems," *Renewable and Sustainable Energy Reviews*, vol. 42, pp. 361–376, 2015.

[24] S. Ray, N. Gupta, and R. A. Gupta, "A comprehensive review on cascaded H-bridge inverter-based large-scale grid-connected photovoltaic," *IETE Technical Review*, vol. 34, no. 5, pp. 463–477, 2017.

[25] P. Sochor and H. Akagi, "Theoretical and experimental comparison between phase-shifted PWM and level-shifted PWM in a modular multilevel SDBC inverter for utility-scale photovoltaic applications," *IEEE Transactions on Industry Applications*, vol. 53, no. 5, pp. 4695–4707, 2017.

[26] ——, "Which is more suitable to a modular multilevel SDBC inverter for utility-scale PV applications, phase-shifted PWM or level-shifted PWM?" in *2016 IEEE Energy Conversion Congress and Exposition (ECCE)*. IEEE, 2016, pp. 1–7.

[27] P. K. Roy, M. Shahjalal, T. Shams, A. Fly, S. Stoyanov, M. Ahsan, and J. Haider, "A critical review on battery aging and state estimation technologies of lithium-ion batteries: prospects and issues," *Electronics*, vol. 12, no. 19, p. 4105, 2023.

[28] A. Vasylyev, A. Vannoni, and A. Sorce, "Optimal dispatch of Li-ion battery energy storage, reviewing and considering cycling and calendar ageing models," *Applied Thermal Engineering*, vol. 265, p. 125597, 2025.

[29] K. K. Duru, C. Karra, P. Venkatachalam, S. A. Betha, A. A. Madhavan, and S. Kalluri, "Critical insights into fast charging techniques for lithium-ion batteries in electric vehicles," *IEEE Transactions on Device and Materials Reliability*, vol. 21, no. 1, pp. 137–152, 2021.

[30] L. Lam, P. Bauer, and E. Kelder, "A practical circuit-based model for Li-ion battery cells in electric vehicle applications," in *2011 IEEE 33rd International Telecommunications Energy Conference (INTELEC)*. IEEE, 2011, pp. 1–9.

[31] C. Zhang, J. Jiang, L. Zhang, S. Liu, L. Wang, and P. C. Loh, "A generalized SOC-OCV model for lithium-ion batteries and the SOC estimation for LNMCO battery," *Energies*, vol. 9, no. 11, p. 900, 2016.

[32] J. Nájera, J. R. Arribas, R. M. de Castro, and C. S. Núñez, "Semi-empirical ageing model for LFP and NMC Li-ion battery chemistries," *Journal of Energy Storage*, vol. 72, p. 108016, 2023.

[33] B. Xu, A. Oudalov, A. Ulbig, G. Andersson, and D. S. Kirschen, "Modeling of lithium-ion battery degradation for cell life assessment," *IEEE Transactions on Smart Grid*, vol. 9, no. 2, pp. 1131–1140, 2018.

[34] P. M. Attia, A. Bills, F. B. Planella, P. Dechent, G. dos Reis, M. Dubarry, P. Gasper, R. Gilchrist, S. Greenbank, D. Howey *et al.*, ""Knees" in lithium-ion battery aging trajectories," *Journal of The Electrochemical Society*, vol. 169, no. 6, p. 060517, 2022.

[35] Q. Ouyang, J. Chen, J. Zheng, and Y. Hong, "SOC estimation-based quasi-sliding mode control for cell balancing in lithium-ion battery packs," *IEEE Transactions on Industrial Electronics*, vol. 65, no. 4, pp. 3427–3436, 2018.

[36] Q. Ouyang, W. Han, C. Zou, G. Xu, and Z. Wang, "Cell balancing control for lithium-ion battery packs: A hierarchical optimal approach," *IEEE Transactions on Industrial Informatics*, vol. 16, no. 8, pp. 5065–5075, 2020.

Shida Jiang received the B.S. degree in electrical engineering from Xi'an Jiaotong University, Xi'an, China, in 2022, and the M.S. degree in systems engineering from the University of California, Berkeley, CA, USA, in 2023. He is currently pursuing the Ph.D. degree in systems engineering with the University of California, Berkeley, CA, USA.

His research interests include battery state estimation, Kalman filtering, and charging optimization.

Shengyu Tao received the Ph.D. degree in electrical engineering through a joint training program with Tsinghua University, Beijing, China, and the University of California, Berkeley, CA, USA. He was a Visiting Research Scholar with the Energy, Controls, and Applications Laboratory (eCAL), University of California, Berkeley. He is currently a Marie Skłodowska-Curie Fellow with the Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden.

His research interests include artificial-intelligence-driven applications for the sustainable utilization of batteries, with particular attention to transferable diagnostics and prognostics under limited and heterogeneous data.

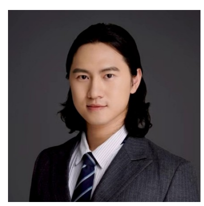

Junzhe Shi received the B.S. degree in mechanical engineering from Purdue University, West Lafayette, IN, USA, in 2016, and the M.S. and Ph.D. degrees in systems engineering from the University of California, Berkeley, CA, USA, in 2018 and 2025, respectively.

He was previously a Battery Algorithm Software Engineer with Apple. He is currently the Co-Founder and Chief Technology Officer of Quotr AI. His research interests include agentic AI, multimodal learning, autonomous decision-making, optimization and control, and intelligent engineering systems.

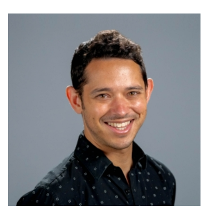

Scott Moura received the B.S. degree from the University of California, Berkeley, CA, USA, in 2006, and the M.S. and Ph.D. degrees from the University of Michigan, Ann Arbor, MI, USA, in 2008 and 2011, respectively, all in mechanical engineering.

From 2011 to 2013, he was a Post-Doctoral Fellow with the Cymer Center for Control Systems and Dynamics, University of California, San Diego, CA, USA. He is currently a Professor of civil and environmental engineering, the Director of the Energy, Controls, and Applications Laboratory (eCAL), and the California PATH Faculty Director with the University of California, Berkeley. His research interests include control, optimization, and AI for batteries, electrified vehicles, and distributed energy resources.

Vincent Molina photograph and biography not available at the time of publication.