From Cell to System: Accelerated hpc Simulations of BESS Aging under Frequency Regulation and Arbitrage use cases

Surya Mitra Ayalasomayajula[a], Michael Starke[b], Srikanth Allu[a,*]
[a]*Computing and Computational Science Directorate*
[b]*Energy Science and Technology Directorate*
Oak Ridge National Laboratory
Oak Ridge, TN 37831, USA
{ayalasomayas, starkemr, allus}@ornl.gov

***Abstract*—Lithium-ion battery energy storage systems (BESS) packs have emerged as a leading solution for grid-scale energy storage, enhancing resiliency and balancing load fluctuations. Yet, experimental characterization of large-format LIB packs—particularly to assess performance and degradation over hundreds of cycles — demands substantial hardware investment and multi-year testing campaigns. In this work, we couple a hierarchical, physics-based modeling framework agnostic to electrode chemistries with high-performance computing to accelerate systems level evaluation by upto two orders of magnitude. Building on the open-source liionpack platform, we implement cell, module, and pack-scale electrochemical models enriched with mechanistic aging mechanisms and deploy them on an HPC cluster to simulate 150–200 kWh systems over 500–1,000 cycles within days. We subject these virtual BESS to both constant-current cycling and realistic grid service profiles spanning frequency regulation, ramp-rate support, and energy arbitrage—and quantify the resulting degradation patterns. Our results reveal that localized cell aging can induce substantial non-uniformity at module and pack levels, with service-specific cycling protocols driving distinct aging modes. This rapid, multiscale modeling approach provides a powerful design-space exploration tool for optimizing electrical architecture, control strategies, and operational schedules to prolong pack lifetime and lower total cost of ownership.**

***Index Terms*—BESS, Physical degradation mechanisms; Cell ageing; Battery pack ageing; PyBaMM; liionpack; high performance computing.**

I. Introduction

Grid-scale lithium-ion battery energy storage systems (LiBESS) are increasingly deployed to enhance grid resilience and balance variable load, with project decisions grounded in site-specific techno-economic analyses tied to intended use cases and operational schedules [1], [2]. In these analyses, battery aging is assumed for specific applications based on previously recorded or historical data, and in order to ensure that the estimated lifetimes are reached, the deployed battery systems employ different operational techniques to improve longevity [3]. Yet as chemistries, pack topologies, and market products evolve, these priors can drift, creating risk for new deployments; moreover, physical prototyping to risk-free designs is capital-intensive and slow [4]. This motivates high-fidelity, computational assessment that support initial analysis of the working design for specific grid applications with optimized performance and aging before field operation.

Current approaches for modeling grid-scale lithium-ion batteries primarily involve equivalent circuit models (ECMs) and phenomenological models. ECMs use simplified circuit analogies to represent battery dynamics, making them computationally efficient for real-time management and control. Phenomenological models, on the other hand, rely on empirical observations and simplified mathematical descriptions to capture battery behavior, particularly aging and degradation phenomena. However, both approaches have shortcomings; ECMs often lack detailed electrochemical insights and can oversimplify degradation processes, reducing accuracy in predicting long-term behavior. Phenomenological models, while more representative of aging mechanisms, require extensive experimental data for calibration and may not generalize well to varied operational conditions or electrode chemistries.

Different grid services impose distinct cycling "signatures" on grid-scale lithium-ion batteries, and those signatures drive different aging pathways. Frequency regulation (FR) usually demands rapid, shallow, bidirectional adjustments—many partial cycles around mid-state-of-charge—whereas energy arbitrage (EA) tends to use fewer but much deeper, energy-intensive cycles aligned with price spreads. Empirical studies consistently show that deep-cycle, high-throughput use accelerates capacity fade relative to shallow FR duty; in controlled testing, FR was the least damaging profile, while

This manuscript has been authored by employee at UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The material is based upon work supported by the U.S. Department of Energy, Office of Electricity (OE), Energy Storage Division.

arbitrage produced the highest degradation for the same service duration or energy processed. These findings reflect the fundamental sensitivity of aging to depth of discharge, C-rate, and temperature, reinforcing that the duty cycle itself—its amplitude, frequency content, and thermal load—materially changes lifetime outcomes.

When FR and EA are stacked, the resulting load profile is a superposition of fast, small-amplitude regulation on top of a slower daily charge/discharge swing. Intuitively this could compound wear, but recent experiments indicate a more nuanced picture: once deep arbitrage cycling dominates electrode fatigue, overlaying shallow FR often adds only modest incremental degradation, even as it can improve revenue and asset utilization—though the incremental impact depends on the relative share of arbitrage energy, SoC operating window, and thermal control. This variability—and the need to co-optimize dispatch, SoC reserves, and cooling while accounting for path-dependent, chemistry-specific aging—motivates a high-fidelity study that couples realistic market signals with physics-informed or validated degradation models, rather than relying on average-cycle abstractions. Such studies are essential to quantify life-throughput trade-offs and design control strategies that let BESS support multiple services while minimizing lifetime cost.

II. Methods

This section describes a grid-scale lithium-ion battery energy storage system (LiBESS) modeling framework. Pack-level simulations were performed with the open-source `liionpack` framework [5], which resolves cell-scale physics by solving a physics-based model for every cell within series/parallel-connected module circuits and updating branch currents at each time step via a circuit solve based on Kirchhoff's laws, *i.e.*, modified nodal analysis [6]. At the cell level, the Single Particle Model with electrolyte (SPMe) [7], [8] implemented in PyBaMM [7] was used; SPMe captures much of the fidelity of a pseudo-2D (P2D) description while remaining computationally efficient [8].

The validated cell and module scale, physics-based parameters from [9] were adopted for large-scale pack analysis. Modules were constructed by connecting eight cells in parallel (1s8p). Packs were then assembled by series/parallel aggregation of modules; for example, connecting 200 modules in series to form a string and then connecting 10 such strings in parallel yields a 200s10p (1s8p) pack. The current simulation framework utilizes the open-source available `liionpack` v0.4.0, and `PyBaMM` v24.9.

The available framework exhibited significant bottlenecks when simulating packs with more than 1000 cells. In the current implementation, these have been addressed to enable simulations with more than 10,000 cells using high-performance computing (HPC) resources. Key improvements include:

- *Ease of hierarchical pack construction*: Easy construction of hierarchical pack circuits using simplified commands, [9], allowing non-uniform resistance elements within a circuit.
- *Memory efficiency*: Efficiently write the outputs to a file periodically, reducing the memory load.
- *Parallel batch solve orchestration on HPC*: Ray-based cluster/actor orchestration for efficient parallel batch solves of cells [10], [11].
- *DFN model for single cells*: optional P2D model for high-fidelity analysis of individual cells within a megapack, useful for mechanistic degradation studies.
- *Numerical solvers*: support for additional solvers, e.g., IDAKLU [12], to increase computational speed and enable future GPU scaling.

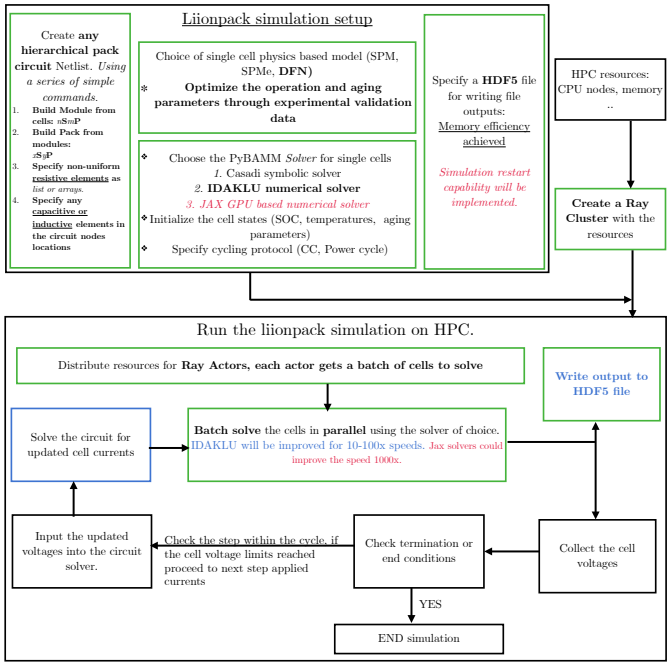

Fig. 1. A schematic of the applied method highlighting the simulation steps and an updated liionpack workflow replicated from [5].

The schematic of the modeling workflow is presented in Figure 1, the features highlighted in green boxes were implemented and working as expected, the blue text and boxes show that particular step works but has a potential to be improved, while text in red highlights future improvements. The current simulation were carried out using the CADES [13] HPC cluster. All the simulations utilized 512 cpu cores distributed across 16 nodes, utilizing 128GB memory. The simulations have a maximum allowed walltime of 2 days, which runs nearly 200 CC cycles of the packs studied with iteration time for each step is nearly 8s.

III. Results and Discussions

Current State-of-the-Art High-voltage BESS systems typically operate above 600V, with industry trends pushing operational voltages upwards to 1000V and beyond, driven by efficiency and cost-reduction targets [14]. Leading battery chemistries for high-voltage applications include Lithium Iron Phosphate (LFP) and Nickel-Manganese-Cobalt (NMC), with

LFP recognized for its robust cycling capabilities under high-voltage conditions [15], [16] and longer cycle life. Further, for a real grid use case, market analysis suggest allowing simultaneous EA and FR services could increase the annual revenue, while still managing degradation [16], [17].

TABLE I
DIFFERENT PACK DESIGNS AND CYCLING CONDITIONS TESTED.

Pack Design	Max Theoretical Capacity (Ah)	Max Theoretical Voltage (V)	Operating Conditions
50s40p (1s8p)	960	210	CC (480A charge 1800s—rest 1800s—480A discharge 1800s—rest 1800s)
100s20p (1s8p)	480	420	CC (240A charge 1800s—rest 1800s—240A discharge 1800s—rest 1800s)
200s10p (1s8p)	240	840	CC (120A charge 1800s—rest 1800s—120A discharge 1800s—rest 1800s)
400s5p (1s8p)	120	1680	CC (60A charge 1800s—rest 1800s—60A discharge 1800s—rest 1800s)
200s10p (1s8p)	240	840	Frequency Regulation (100%)
200s10p (1s8p)	240	840	Energy Arbitrage (100%)
200s10p (1s8p)	240	840	Combination of x% FR + y% EA

In this regard, different pack configurations ranging from low voltage to high voltage systems with same energy rating of $\sim$ 175kWh, are tested with CC cycling which represent nearly the same power. To showcase potential reasons for migrating to high-voltage systems. Further, the capability of the framework to simulate grid use cases are highlighted. Table I, highlights the pack design and operating conditions.

Fig. 2 highlights the effects of Li BESS design on the cell-modules-packs hierarchical aging predictability. Fig. 2a highlights the percentage capacity loss at the pack scale; Fig. 2b shows the percentage capacity loss at the module scale; Fig. 2c shows the percentage capacity loss at the individual cell scale. The aging results show that, despite the battery packs aging similarly under the same cycling conditions, the high-voltage systems have lower aging variation across the cells and modules, due to fewer parallel connections causing lower current imbalances, while a low-voltage system has the most aging variation across individual cells and modules. Further, in Fig. 2a the 50s40p (1s8p) low-voltage system shows less degradation at the pack scale, since the simulations did not allow overcharging of cells/modules, which made the simulation proceed to the next step (rest) after the cells hit the high-voltage limit, as seen in Fig. 2d. This means that in a real system, monitoring aging only at the pack scale could be unreliable, as overcharging can cause increased aging of cells and catastrophic failure. Nevertheless, the results show that high-voltage systems age reliably in comparison to a low-voltage system highlighting the industrial shift towards high-voltage systems.

Fig. 3 shows the statistical metrics for aging of individual cells in different battery pack configurations. The pack capacity loss percentage matches closely with the mean of the individual cell capacity percentage, since cell/module capacities are additive in parallel connections and the cell/modules with minimum capacity dictates aging of serially connected cells/modules. The variance bands show the standard deviation in aging predictions, the low-voltage systems have larger variations since, the number of parallel connections determine the level of current imbalances experienced, while high-voltage systems have smaller variance bands implying reliable aging predictions. The lower median relative to the mean signals a right-skewed aging distribution i.e., most cells show limited degradation, but a few age significantly, elevating module/pack failure risk. The mean–median gap thus provides a simple

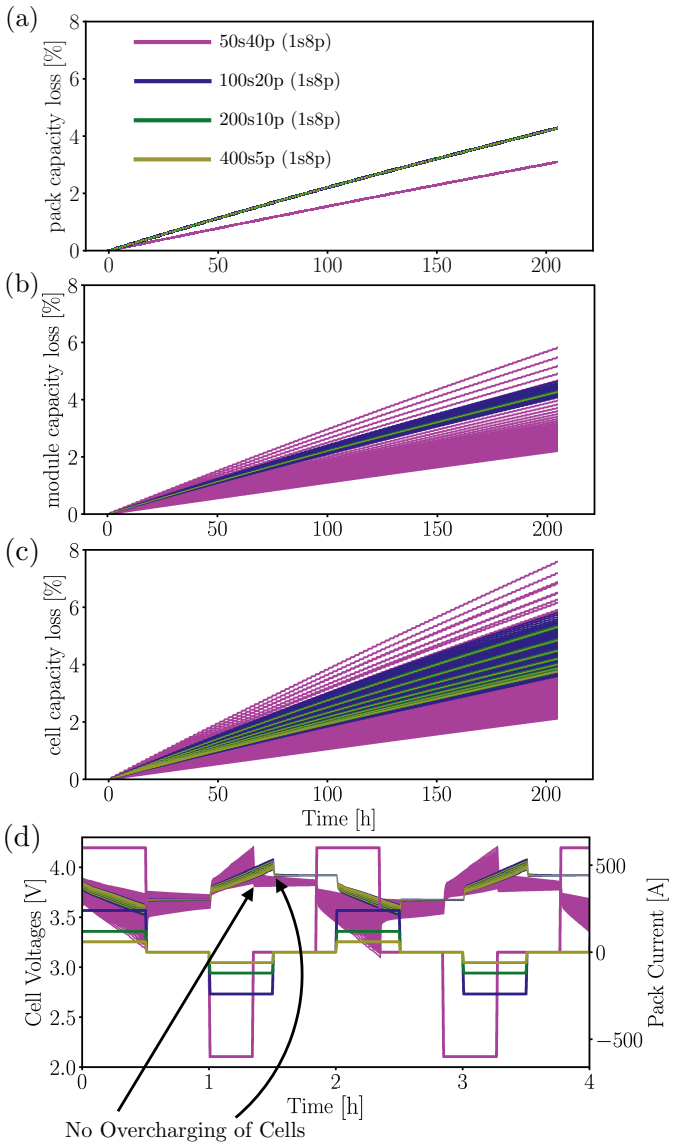

Fig. 2. Constant Current cycling of four different 175kWh LiBESS system configurations. Insets (a) highlight the percentage loss in capacity of the packs, (b) modules within the pack, and (c) individual cells. Inset (d) shows individual cell voltages, and the pack currents. The results highlight the high-voltage systems degrade reliably since the individual cell level aging heterogeneity minimizes.

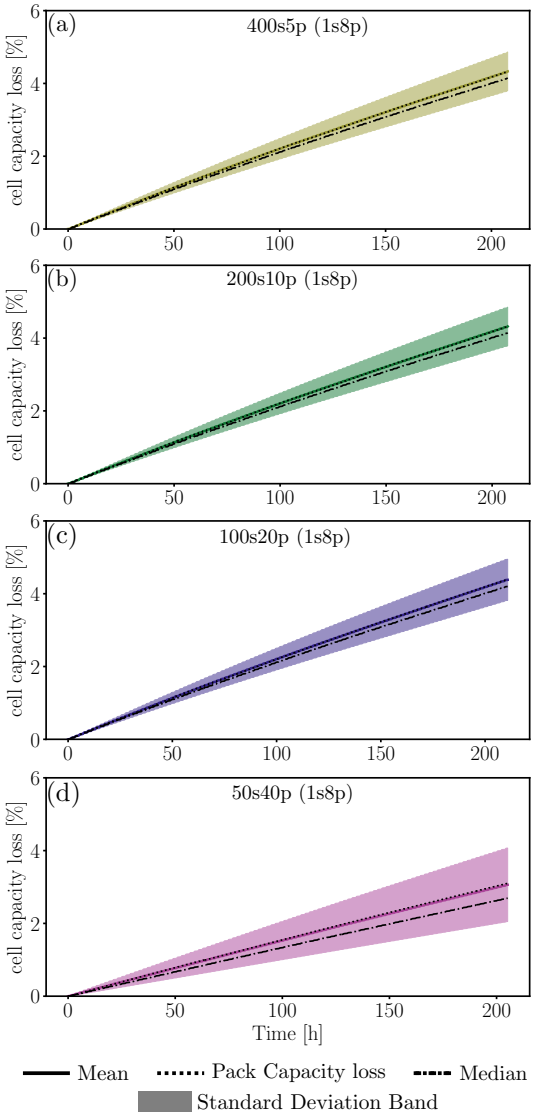

Fig. 3. Statistical metrics: mean, median and standard deviation, for individual cell aging in different pack design configurations.

indicator of heterogeneity and prediction reliability.

Fig. 4 illustrates common LiBESS grid-service use cases. The 24 h power cycles for energy arbitrage (EA), frequency regulation (FR), and their combinations are considered. The normalized FR profile is taken from Fig. 5.3.2 of [18], while the EA profile is arbitrarily chosen based on the literature [16]. The energy within the cycle is defined as $E_{\text{cycle}} = \int_0^{24\,\text{h}} P_{\text{cycle}}\,\mathrm{d}t$ and is constrained to zero over 24 h so the pack returns to the same state of charge (SOC) without explicit rebalancing [18]. For mixed use cases with $x\%$ EA and $y\%$ FR, the 24 h profile is constructed so that the ratio of energy contributions from EA and FR is $x : y$.

The applied power is $P_{\text{cycle}} = P_{\text{Norm}} \times P_{\text{N}}$, where P_{Norm} is the normalized profile and P_{N} is the nominal pack power. Pack capacity $Q_{\text{pack}} = Q_{\text{rated,cell}} \times N_{\text{P,mod}} \times N_{\text{P,pack}}$, with $Q_{\text{rated,cell}}$, the rated cell capacity, $N_{\text{P,mod}}$, cells in parallel per module, and $N_{\text{P,pack}}$, modules in parallel per pack. The pack voltage

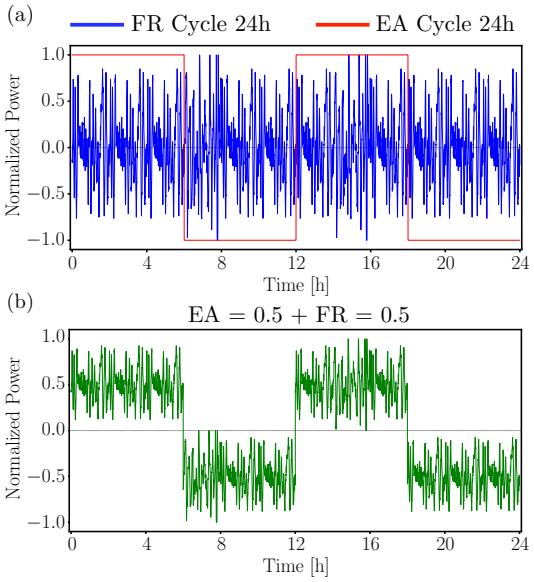

Fig. 4. Normalized power profiles for generalized grid use cases such as energy arbitrage EA and frequency regulation FR.

is $V_{\text{pack}} \approx V_{\text{cell,Nom}} \times N_{\text{S,mod}} \times N_{\text{S,pack}}$, with $V_{\text{cell,Nom}}$, the nominal cell voltage, $N_{\text{S,mod}}$, series cells per module, and $N_{\text{S,pack}}$, series modules per pack. Hence, $E_{\text{pack}} = Q_{\text{pack}} \times V_{\text{cell,Nom}}\ times N_{\text{S,mod}} \times N_{\text{S,pack}}$, where E_{pack}, is the nominal pack energy.

For a 200s10p (1s8p) configuration, $(E_{\text{pack}})_{200\text{s}10\text{p}(1\text{s}8\text{p})} = 240 \times 3.6 \times 200 \times 1 = 172.48$ kWh. Assuming the pack draws/delivers the energy at 1C, the nominal power is $P_{\text{N}} = 172.48$ kW.

In Figure 5, expected trends for different grid use cases are observed, the pure EA cycle (EA = 1.0) produces the largest and fastest pack capacity loss, while the pure FR cycle (FR = 1.0) ages the least, with mixed EA/FR cases falling in between. Pure EA cycle causes most aging as expected due to a deeper charge and discharge, in comparison to a pure FR cycle which operates within a narrow range of SOC [16]. Mechanistically, sustained FR can still drive particle-scale electrochemo-mechanical fatigue: oscillatory power induces repeated intercalation strains that promote microcracking and contact loss. By contrast, EA's deeper SOC excursions are more likely to accentuate SEI growth and, under unfavorable conditions (e.g., low temperature or high rates), lithium plating, leading to larger irreversible capacity loss. Therefore, from an operational standpoint, the framework is capable to identify operational strategies, i.e., a mixed EA/FR schedules offer a tunable trade-off between revenue and degradation and could be co-optimized with maintenance intervals and warranty targets.

IV. Summary and Conclusions

The existing liionpack modeling framework has been improved to incorporate features that facilitate physics based hierarchical analysis of kWh–GWh grid-scale LiBESS systems for various operating conditions. The results show that the characteristic aging observed at the system scale has a

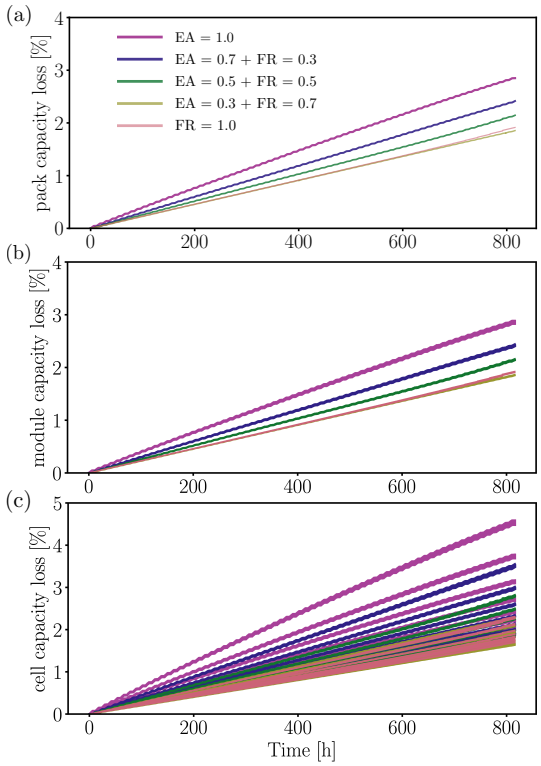

Fig. 5. Hierarchical aging of a 200s10p (1s8p) battery pack under grid usecase: EA, FR and combined EA+FR. Insets show capacity loss percentage at, (a) pack scale, (b) module and (c) individual cells. EA cycle shows more degradation in comparison to the FR, while all other EA/FR combinations tend to fall in between.

variation with respect to individual cells depending on the pack design. Therefore, the predictability of long-term aging of the system is expected to vary largely due to few cells hitting the end of life faster.

Key findings. (i) *Pack Design Matters:* high-voltage designs (more series, fewer parallel strings) exhibit markedly lower cell-to-cell and module-to-module aging dispersion than low-voltage designs; pack-level metrics alone can mask emerging risk, highlighting the need to resolve cell/module scales for better analysis. (ii) *Duty-cycle tuning:* pure EA induces the most capacity loss; pure FR ages the least; mixed EA/FR falls in between, thus allow tuning of EA/FR ratios to balance revenue and aging. (iii) *Statistical indicators:* right-skewed cell-aging distributions (mean $>$ median) highlight heterogeneity; the mean–median gap and standard deviation provide compact reliability metrics.

Although the high-voltage systems look promising, these pose maintenance difficulties and the formation of thermal hot spots will be analyzed in the future. The framework could be used to study non-zero daily energy budgets, broader chemistries, and field conditions. *Throughput*– with the current 2-day wall time we complete $\sim$ 200 cycles; extrapolating, 500–1000 cycles are achievable in $\sim$ 6–10 days. By parallelizing the circuit solve and file write operations, we target 1000 cycles within 2 days (reducing per-step time from $\sim$8 s to $\sim$ 0.5 s). Further, the framework will leverage optional DFN/P2D models for cells, incorporate explicit fatigue/plating submodels, and exploit numerical solvers (and GPU based solvers) to further improve computing efficiency by $\sim$10x-100x, and expand design-space exploration for pack design topology, operational control, and optimize dispatch for grid use cases.

References

[1] C.-T. Tsai, E. M. Ocampo, T. M. Beza, and C.-C. Kuo, "Techno-economic and sizing analysis of battery energy storage system for behind-the-meter application," *IEEE Access*, vol. 8, pp. 203 734–203 746, 2020.

[2] P.-H. Hsi and J. C. P. Shieh, "Techno-economic investment risk modeling of battery energy storage system participating in day-ahead frequency regulation market," *IEEE Access*, vol. 12, pp. 56 981–56 990, 2024.

[3] R. Colucci, I. Mahgoub, H. Yousefizadeh, and H. Al-Najada, "Survey of strategies to optimize battery operation to minimize the electricity cost in a microgrid with renewable energy sources and electric vehicles," *IEEE Access*, vol. 12, pp. 8246–8261, 2024.

[4] E. Hanson, S. Mayekar, and V. Dravid, "Applying insights from the pharma innovation model to battery commercialization—pros, cons, and pitfalls," *MRS Energy & Sustainability*, vol. 4, 2017.

[5] T. Tranter, R. Timms, V. Sulzer, F. Planella, G. Wiggins, S. Karra, P. Agarwal, S. Chopra, S. Allu, P. Shearing *et al.*, "liionpack: A python package for simulating packs of batteries with pybamm," *Journal of Open Source Software*, vol. 7, no. 70, 2022.

[6] C.-W. Ho, A. Ruehli, and P. Brennan, "The modified nodal approach to network analysis," *IEEE Transactions on circuits and systems*, vol. 22, no. 6, pp. 504–509, 1975.

[7] V. Sulzer, S. G. Marquis, R. Timms, M. Robinson, and S. J. Chapman, "Python battery mathematical modelling (pybamm)," *Journal of Open Research Software*, vol. 9, no. 1, 2021.

[8] S. G. Marquis, V. Sulzer, R. Timms, C. P. Please, and S. J. Chapman, "An asymptotic derivation of a single particle model with electrolyte," *Journal of The Electrochemical Society*, vol. 166, no. 15, p. A3693, 2019.

[9] S. M. Ayalasomayajula, Y. Preger, J. Mueller, and S. Allu, "Physics-based analysis of cell imbalances and aging in lithium-ion battery modules and packs," *Journal of The Electrochemical Society*, 2025.

[10] Ray Project, "Ray," https://github.com/ray-project/ray, 2025, apache-2.0 license; latest release v2.48.0 (July 18, 2025).

[11] P. Moritz, R. Nishihara, S. Wang, A. Tumanov, R. Liaw, E. Liang, M. Elibol, Z. Yang, W. Paul, M. I. Jordan, and I. Stoica, "Ray: A distributed framework for emerging AI applications," in *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*. Carlsbad, CA: USENIX Association, Oct. 2018, pp. 561–577. [Online]. Available: https://www.usenix.org/conference/osdi18/presentation/moritz

[12] PyBaMM Team, "Idaklu solver — PyBaMM manual," https://docs.pybamm.org/en/stable/source/api/solvers/idaklu_solver.html, 2025, pyBaMM v25.8.0 (API docs).

[13] Oak Ridge National Laboratory (CADES). (2025) Cades user documentation. Compute and Data Environment for Science (CADES) user documentation. [Online]. Available: https://docs.cades.ornl.gov/

[14] J. He, Y. Yang, and D. Vinnikov, "Energy storage for 1500 v photovoltaic systems: A comparative reliability analysis of dc-and ac-coupling," *Energies*, vol. 13, no. 13, p. 3355, 2020.

[15] J. Olmos, I. Gandiaga, A. Saez-de Ibarra, X. Larrea, T. Nieva, and I. Aizpuru, "Modelling the cycling degradation of li-ion batteries: Chemistry influenced stress factors," *Journal of Energy Storage*, vol. 40, p. 102765, 2021.

[16] B. Ellis, C. White, and L. Swan, "Degradation of lithium-ion batteries that are simultaneously servicing energy arbitrage and frequency regulation markets," *Journal of Energy Storage*, vol. 66, p. 107409, 2023.

[17] N. B. Arias, J. C. López, S. Hashemi, J. F. Franco, and M. J. Rider, "Multi-objective sizing of battery energy storage systems for stackable grid applications," *IEEE Transactions on Smart Grid*, vol. 12, no. 3, pp. 2708–2721, 2020.

[18] D. R. Conover, A. J. Crawford, V. V. Viswanathan, S. Ferreira, and D. Schoenwald, "Protocol for uniformly measuring and expressing the performance of energy storage systems," Pacific Northwest National Lab.(PNNL), Richland, WA (United States), Tech. Rep., 2014.