Degradation-aware sizing and scheduling of a solar-integrated two-wheeler battery-swapping station with heterogeneous pack health

A Monte Carlo electro-thermal ageing study for an Indian urban kiosk

[1]Pravat Satpathy, [2]Sonia Mohapatra, [3]Saswat Swain
[1]Researcher, [2]Researcher, [3]Ph.D Scholar
[1,2]Research and Development Division,
Tekstein Solutions Private Limited, Bhubaneswar, Odisha, India
[3]KSOM — KIIT Deemed to be University,
Bhubaneswar, Odisha, India
[1]analytics.pravat@gmail.com, [2]smctc98@gmail.com, [3]saswat.mck@gmail.com

***Abstract*—Battery swapping can make electric two-wheelers economical for high-utilisation gig work in India, and rooftop photovoltaics can cut the cost of the energy swapped. Existing sizing models for solar-integrated swap stations, however, treat every battery pack as identical, although shared packs age at different rates and a weak pack fails a long-distance rider while still serving a commuter. This paper couples station sizing with per-pack health. A closed-loop 5-min simulation of a Bhubaneswar kiosk gives each lithium iron phosphate pack its own Arrhenius cycle and calendar fade, resistance growth and temperature, and serves 30 riders in three duty classes under calibrated synthetic weather. A degradation-aware swap scheduling policy hands the healthiest pack to high-need riders and the weakest adequate pack to low-need riders on a 95th-percentile need forecast, charges from solar unless a least-laxity test requires grid energy, and lets the sizing loop choose the retirement depth. Designs are selected under worst-year service limits using 4,610 five-year Monte Carlo runs in screening, verification and independent reporting stages. An identical-battery planning model understated equivalent annual cost by 15% and worst-year mid-trip swaps by about 3 times. Health-aware dispatch reduced worst-year mid-trip swaps by 14–25% at negligible direct cost, which made deep retirement feasible: the best health-aware design met the service limit with 2 fewer slots and 5 kWp less photovoltaics, at 3.3% lower cost than the best health-blind design. The benefit rose to 9.3% under a tighter limit and vanished under a loose one, while wear-levelling dispatch never helped. Network-level rebalancing and hardware validation are left for future work.**

***Index Terms*—Battery swapping, electric two-wheelers, state of health, lithium iron phosphate, battery degradation, rooftop photovoltaics, station sizing, Monte Carlo simulation, dispatch scheduling**

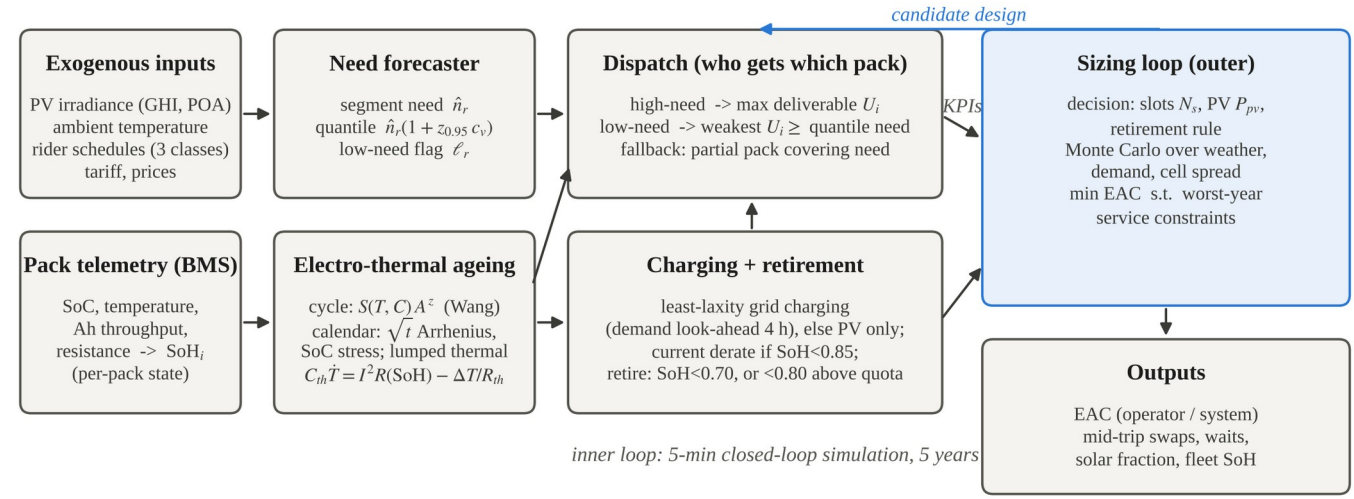


Fig. 1. Degradation-aware swap scheduling (DASS) framework. The inner loop simulates the kiosk every 5 min for five years with per-pack electro-thermal ageing; the outer loop searches the number of slots, the PV rating and the retirement rule under worst-year service constraints.

I. Introduction

Electric two-wheelers (E2Ws) carry most of India's urban gig mobility, and battery swapping removes the longest friction in that business: the hour or more a rider would otherwise lose at a charger. The economic case is narrower than the enthusiasm suggests. A total-cost-of-ownership (TCO) analysis by the International Council on Clean Transportation (ICCT) found that swapping beat point charging only for ride-hailing duty, where the Okinawa iPraise+ with swapping came out 15% cheaper per kilometre, while personal and last-mile delivery use remained cheaper on point charging [1]. The follow-up ICCT/NITI Aayog report located the break-even near 140 km/day and identified the energy price per swapped kWh, not hardware, as the dominant cost lever [2]. Swapping therefore pays when two things are true at once: riders drive far enough that their time is worth protecting, and the operator delivers each kWh cheaply.

Rooftop or canopy photovoltaics (PV) attack the second condition directly, and the stock of charged packs at a kiosk is, in effect, a distributed energy store that can absorb midday solar output. Manoj *et al.* sized a rooftop-PV E2W swap station for an Indian residential community with Monte Carlo treatment of arrival and charging uncertainty [3], and related work sized PV-based swap stations at petroleum retail outlets with linear programming [4]. These studies share a structural simplification with most of the swap-station literature [5–7]: every pack is the same pack. Batteries enter as a pooled inventory of identical units, and ageing, if represented at all, is a linear cost per kWh of throughput.

That simplification hides the variable that decides whether the station keeps its service promise. Swappable packs are shared by many riders, so each pack's state of health (SoH) is the record of everybody's usage, and manufacturing spread makes nominally identical packs age at visibly different rates [8], [9]. The ICCT/NITI report names the consequence plainly: a swap subscriber has no assurance of consistent range from whichever pack the kiosk hands over [2]. A pack at 78% SoH is an adequate battery for a commuter who drives 30 km a day and an inadequate one for a ride-hailing rider who needs 35 km between planned stops. Once packs differ, *which pack goes to which rider* becomes a decision with a cost, and the station's sizing (how many packs and slots, how much PV, when to retire a pack) can no longer be separated from that dispatch decision.

This paper makes that coupling explicit. Each pack carries its own electro-thermal ageing state, riders carry stochastic segment-level energy needs, and the operator's dispatch, charging and retirement rules are evaluated jointly with the station size over a five-year Monte Carlo horizon. The research questions are: (Q1) how wrong is an identical-battery planning model once heterogeneous ageing is simulated; (Q2) which part of a health-aware policy—dispatch, charging or retirement—creates value, and which does not; and (Q3) how does the cost-optimal, service-feasible station size change when sizing and degradation-aware scheduling are solved together. The main contributions are:

- A closed-loop, 5-min resolution model of a PV-integrated E2W swap station in which every pack has its own capacity, resistance, temperature and ageing trajectory (Arrhenius cycle and calendar fade with pack-to-pack rate dispersion), coupled to rider work segments of three duty classes and a calibrated synthetic weather year for Bhubaneswar, India.
- A degradation-aware swap scheduling (DASS) policy with three parts: need-stratified assortative dispatch on quantile need forecasts, least-laxity solar-aware charging, and a retirement depth chosen inside the sizing loop (fixed thresholds or a demand-matched quota).
- Two analytical results that explain the simulation: wear-levelling dispatch *maximises* total fleet capacity fade under a concave fade law (Proposition 1), and assortative matching is near-optimal only when the weak pack still covers the low-need rider's high-quantile demand (Proposition 2), which is why naive best-fit dispatch fails.
- A service-constrained sizing study over 4,610 five-year simulations organised as screening, verification and independent reporting stages with common random numbers, a full factorial attribution, and explicit reporting of negative results and of two corrections made during the study.

Section II reviews related work, Section III states the problem, Sections IV and V present the model and the DASS method, Section VI describes the simulation design, Section VII reports results, Section VIII compares with prior work, and Sections IX–X summarise outcomes, limitations and future work.

II. Related Work

Table I summarises the studies closest to this work. Operational scheduling of swap stations was formalised as a mixed-integer program for day-ahead charging and battery-to-grid services, with degradation linearised in charge and discharge power [5]. Planning studies compared swap and charging stations in distribution systems [6]. For renewable-powered stations, Renga *et al.* traded quality of service against PV size and smart charging for urban scooter stations [7], and Rohde *et al.* designed an off-grid PV swap kiosk for low-income operators [10]. The Indian studies [3], [4] are the direct antecedents: both size PV against stochastic demand, and both treat the pack pool as homogeneous.

Battery-specific degradation has entered swap-station *operation* only recently. Chen *et al.* derived an intertemporal marginal degradation cost to coordinate charging, swapping and grid services [11]; an *IEEE Transactions on Smart Grid* study addressed station management under heterogeneous battery degradation [12]; Shalaby *et al.* trained a deep reinforcement learning agent with degradation and waiting penalties [13]; and Li *et al.* embedded a physics-based ageing surrogate in model predictive control, choosing which batteries to release while serving requests first-in-first-out [14]. None of these couples per-pack health with station sizing and PV, none conditions the hand-out on the receiving rider's energy need, and the MPC study explicitly leaves vehicle-level dispatch to FIFO [14]. The gap addressed here is therefore the joint problem: per-pack health, rider-specific need, solar-coupled charging and retirement inside the sizing loop.

TABLE I
COMPARISON OF CLOSELY RELATED STUDIES

Ref.	Method	Per-pack SoH	PV	Sizing	Rider-aware dispatch	Main limitation
[5]	MILP, day-ahead	No (linear)	No	No	No	Identical packs
[6]	Planning, LCC	No	No	Yes	No	No ageing physics
[7]	Sim., QoS–PV trade-off	No	Yes	Yes	No	Identical packs
[3]	Monte Carlo sizing	No	Yes	Yes	No	Identical packs
[4]	LP (COIN-OR)	No	Yes	Yes	No	Deterministic, identical
[11]	Intertemporal MDC	Pooled	No	No	No	No sizing, no PV
[13]	Deep RL	Penalty	No	No	No	Opaque, no sizing
[14]	MPC + SPM surrogate	Yes	No	No	No (FIFO)	No PV, fixed size
This work	MC sim. + DASS	Yes	Yes	Yes	Yes	Simulation only

III. PROBLEM STATEMENT

Consider a single kiosk with N_s charging slots serving M subscribed riders, each of whom holds exactly one pack, so the closed fleet has $N = N_s + M$ packs. A canopy PV array of rating P_{pv} and a grid connection supply the slots. The operator chooses (i) the design $x = (N_s, P_{pv})$, and (ii) an operating policy π that decides, at every swap request, which pack to hand over, at every 5-min step, which packs to charge from which source, and, at every deposit, whether to retire the returned pack. The engineering difficulties are fourfold. First, packs are heterogeneous and their capacities drift apart over time, so the deliverable energy of the inventory is not N_s times a constant. Second, rider needs are stochastic and differ by duty class, so a pack adequate for one rider fails another. Third, PV output peaks when ambient temperature peaks and when few riders are at the kiosk, so solar-first charging competes with readiness and with thermal ageing. Fourth, retiring packs late saves capital but erodes service, and the erosion is invisible to a model that never lets packs age. The practical constraints are an unchanged 2-kWh pack format, a flat EV tariff with no demand charge [15], no export credit for surplus PV (relaxed in sensitivity), and a service level the operator must keep in *every* year of the horizon, not only on average.

IV. SYSTEM MODEL

A. Station, PV and Weather

The site is Bhubaneswar, Odisha (20.25° N, 85.83° E). Clear-sky extraterrestrial irradiance follows standard solar geometry [16]. Daily clearness indices are drawn from a Beta distribution around a monthly mean with an AR(1) Gaussian copula (lag-one correlation 0.55) to reproduce cloudy spells; an AR(1) intra-day term adds 5-min fluctuation. Beam and diffuse parts follow the Erbs correlation [17] and are transposed to a 20° south-facing plane with an isotropic sky. Array output is

$$P(t) = P_{pv}\,(G_{poa}/1000)\,[1 + \beta(T_c - 25)]\,\eta_d \qquad (1)$$

with cell temperature $T_c = T_a + G_{poa}(T_{NOCT} - 20)/800$, where G_{poa} is plane-of-array irradiance (W/m^2), $\beta = -0.0037$ K^{-1}, $T_{NOCT} = 45$ °C and $\eta_d = 0.86$ lumps inverter, wiring, soiling and mismatch losses. The monthly clearness shape is scaled so that the long-run global horizontal irradiance (GHI) equals the Global Solar Atlas value of 4.87 kWh/m^2/day [18]. Ambient temperature combines the JMA monthly normals [19], a diurnal sinusoid peaking at 14:30 and an AR(1) day anomaly. As an independent check, the model's specific yield is 1438 kWh/kWp/yr against the atlas value of 1477 kWh/kWp/yr (−2.6%), and the simulated monthly mean temperatures lie within 0.3 K of the normals (Fig. 2).

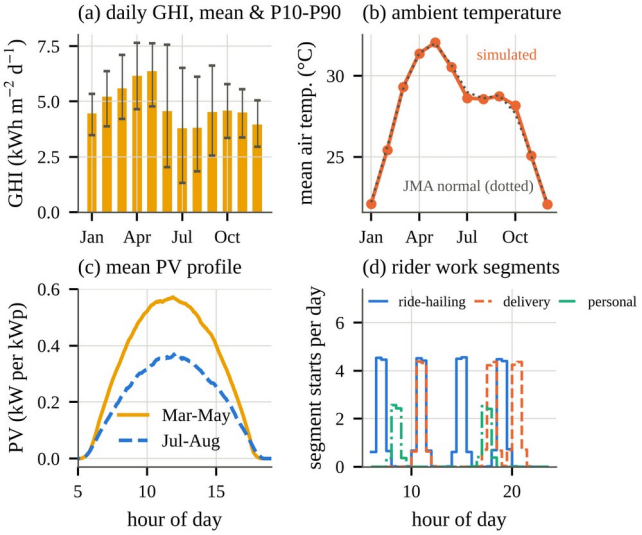


Fig. 2. Model inputs for Bhubaneswar. (a) Daily GHI by month with P10–P90 band. (b) Simulated monthly mean temperature against JMA normals. (c) Mean PV profile in the pre-monsoon and monsoon seasons. (d) Start times of rider work segments by duty class (first simulated year).

B. Rider Demand

Riders belong to three classes chosen to match the ICCT utilisation bands [1]: 12 ride-hailing riders (140 ± 20 km/day, 6 days/week), 10 delivery riders (100 ± 15 km/day, 7 days/week) and 8 personal commuters (30 ± 6 km/day, weekdays). Each rider carries a persistent habit multiplier (log-normal, $\sigma = 0.08$), and energy intensity is $e = 0.032$ kWh/km, consistent with the 111-km range of the reference scooter in [1]. A working day is split into segments: four 3.5-h blocks for ride-hailing, a lunch block and two evening blocks for delivery, and a morning and evening commute for personal use (Fig. 2(d)). The need of rider r in segment k on day d is

$$n_{r,k,d} = s_k \cdot D_{r,d} \cdot e, \quad D_{r,d} \sim \mathrm{N}(m_r, \sigma_c^2) \qquad (2)$$

with s_k the segment share of daily distance $D_{r,d}$. Riders plan their swap points so that a fresh pack covers a segment with margin, which is the rational behaviour when the advertised range of a 2-kWh pack is about 54 km. At each segment start a rider swaps if the remaining deliverable energy is below 1.1 times the planned need (personal riders check the whole day's need, because they pass the kiosk once in the morning). A pack that

runs down to the 10% reserve inside a segment forces an *unplanned* mid-segment swap. Planned swaps cost 5 min of rider time and unplanned ones 20 min (detour plus a lost trip), both valued at INR 65/h, the ICCT value-of-time baseline [2]. If no fully charged pack is available, the rider waits in a FIFO queue for up to 20 min and then leaves with the best partially charged pack.

C. Pack Electro-Thermal Ageing

Packs are 16S LFP modules of 40 Ah (E_{nom} = 2.048 kWh). Pack i has state of health

$$\text{SoH}_i(t) = c_{0,i} - \kappa_i\,[q_i^{\text{cyc}}(t) + q_i^{\text{cal}}(t)]/100 \tag{3}$$

where $c_{0,i}$ N(1, 0.01^2) is the as-built capacity and κ_i logN(0, 0.15^2) is a pack-specific ageing-rate multiplier representing the production-caused dispersion of fade trends [8], [9]. Cycle fade follows the semi-empirical graphite/LFP law of Wang *et al.* [20],

$$q^{\text{cyc}} = B\exp[-(E_a - \alpha C)/(RT)]\,A^z \tag{4}$$

with B = 31630, E_a = 31.7 kJ/mol, α = 370.3 J/mol per C-rate, z = 0.55 and A the reference-cell Ah throughput (half-weight on charge and half on discharge, each evaluated at its own temperature and C-rate, the C-rate being referred to the *current* capacity). Because the stress factor varies over time, (4) is integrated in its equivalent-throughput form, $A_{\text{eq}} = (q^{\text{cyc}}/S)^{1/z}$, $q^{\text{cyc}} \leftarrow S(A_{\text{eq}} + \Delta A)^z$. Calendar fade uses the square-root-of-time, Arrhenius and state-of-charge (SoC) dependence reported for commercial LFP cells [21],

$$q^{\text{cal}} = k_{\text{ref}}\exp[-(E_c/R)(1/T - 1/298.15)]\,(0.6 + 0.8\sigma)\,\sqrt{t} \tag{5}$$

calibrated to 2.0% loss after one year at 25 °C and 50% SoC, with E_c = 40 kJ/mol; σ is SoC. Internal resistance grows as $R = R_0[1 + \gamma(1 - \text{SoH})]$ with R_0 = 30 mΩ and γ = 2. Pack temperature obeys the lumped balance

$$C_{\text{th}}\,\mathrm{d}T/\mathrm{d}t = I^2R(\text{SoH}) - (T - T_{\text{env}})/R_{\text{th}} \tag{6}$$

with C_{th} = 12 kJ/K, R_{th} = 0.6 K/W in a slot ($T_{\text{env}} = T_a$ + 2 K) and 0.8 K/W under the seat ($T_{\text{env}} = T_a$ + 3 K). Under these parameters a pack cycled at 0.5C reaches 20% cycle fade after about 7588 equivalent full cycles (EFC) at 25 °C but only 3584 EFC at 35 °C, so the Bhubaneswar climate, not the cell chemistry, sets the effective life. The deliverable energy of a pack at SoC σ_i is

$$U_i = (\sigma_i - \sigma_{\text{res}})\,\text{SoH}_i\,E_{\text{nom}}, \quad \sigma_{\text{res}} = 0.10 \tag{7}$$

D. Cost and Service Metrics

Over the horizon H = 5 years and real discount rate r = 10%, the equivalent annual cost (EAC) is

$$\text{EAC} = C_{\text{bat}} + C_{\text{slot}} + C_{\text{pv}} + C_{\text{en}} + C_{\text{time}} \tag{8}$$

with the five terms

$$C_{\text{bat}} = \Phi_H[N\,c_p + \Sigma_j\,c_p(1+r)^{-t_{-}\{j\}} - V_{\text{res}}] \tag{9}$$

$$C_{\text{slot}} = N_s c_s(\Phi_{10} + o_s),\;\; C_{\text{pv}} = P_{\text{pv}}c_{\text{pv}}(\Phi_{25} + o_{\text{pv}}) \tag{10}$$

$$C_{\text{en}} + C_{\text{time}} = \Phi_H\,\Sigma_d(\pi_e E_d^g + v\,L_d)(1+r)^{-d/365} \tag{11}$$

where Φ_n is the capital recovery factor over n years, c_p = INR 13,500/kWh × E_{nom} is the pack price, t_j are replacement times, c_s = INR 25,000 per slot, c_{pv} = INR 45,000/kWp, o are O&M fractions, π_e = INR 5/kWh is the Odisha EV-station tariff [15], E_d^g is grid energy and L_d the rider time lost on day d, valued at v = INR 65/h. Residual pack value at H is depreciated linearly to the retirement threshold θ: $V_{\text{res}} = \Sigma_i\,c_p\,\max[0, (\text{SoH}_i - \theta)/(1 - \theta)]\,(1+r)^{-H}$. The *operator* EAC excludes the rider-time term; the *system* EAC includes it, in line with the ICCT treatment of value of time [1]. Service is measured by unplanned swaps per 1000 swaps (u), partial-pack hand-outs per 1000 swaps (p) and mean wait per swap (w).

E. Sizing Problem

With ω denoting the random scenario (weather, rider draws, pack spread) and y the simulation year, the degradation-aware sizing problem is

$$\min_{x,\,\pi\,\in\,\Pi}\,\mathrm{E}_\omega[\text{EAC}(x, \pi, \omega)] \tag{12}$$

$$\max_y \mathrm{E}[u_y] \le 40,\;\; \max_y \mathrm{E}[p_y] \le 10,\;\; \mathrm{E}[w] \le 1 \tag{13}$$

The worst-year form of (13) matters: an ageing fleet degrades service in years 4–5, and an average over the horizon would conceal that. The conventional *identical-battery planning model* used for comparison is the same simulator with ageing switched off (SoH ≡ 1, $\kappa_i \equiv 1$) and pack cost annualised over a datasheet life of 2000 EFC, mirroring the homogeneous treatment in [3], [4].

V. Degradation-Aware Swap Scheduling (DASS)

Fig. 1 shows the framework. The inner loop is the 5-min closed-loop simulation; the outer loop searches designs. DASS needs only quantities a connected kiosk already logs: battery-management-system (BMS) telemetry for SoC, temperature and throughput, an SoH estimate, and each subscriber's history of segment needs [2].

A. Quantile Need Forecast

For rider r the kiosk keeps the mean planned need $\hat{n}_r$ and the class coefficient of variation $c_{v,r}$, and forms the high-quantile need

$$\tilde{n}_r = \hat{n}_r(1 + z_q c_{v,r}),\quad z_q = 1.645\ (95\text{th percentile}) \tag{14}$$

and a low-need flag $\ell_r = 1$ if $\hat{n}_r < 0.65 \cdot 0.85\,E_{\text{nom}}$, i.e., if the rider's typical need is below 65% of a fresh pack's usable energy. In the base case this flags the personal commuters; the rule itself never uses class labels.

B. Need-Stratified Assortative Dispatch

Let **R** be the set of fully charged packs at the kiosk ($\sigma_i \ge 0.95$). For a high-need rider ($\ell_r = 0$) DASS returns the pack with the largest deliverable energy; for a low-need rider it returns the *weakest* full pack that still covers the quantile need:

$$i^* = \text{argmax}_{i\in\mathbf{R}}\,U_i, \quad \text{if } \ell_r = 0 \tag{15}$$

$$i^* = \text{argmin}_{i\in\mathbf{R},\ U_\{i\,\ge\,\tilde{n}_r\}}\,U_i, \quad \text{if } \ell_r = 1 \tag{16}$$

If no full pack covers $\tilde{n}_r$, the largest full pack is issued; if **R** is empty, a partially charged pack with $U_i \ge \tilde{n}_r$ is issued instead of making the rider wait. The baseline it replaces is the industry FIFO rule (the longest-ready full pack). The rule is deliberately *not* best-fit for everyone: Proposition 2 below explains why best-fit on the *mean* need fails, and Section VII quantifies the cost of applying best-fit to every rider.

C. Least-Laxity Solar-Aware Charging

At step t the kiosk lists the instants τ_k within a 4-h look-ahead at which a fresh pack will be demanded: queued riders, riders whose current pack will deplete mid-segment, and riders whose next segment will trigger a swap, plus a buffer of two packs. With R_0 packs already ready and the remaining packs sorted by their time-to-ready $\rho_{(j)}$ at full slot power, the laxity of the k-th uncovered demand is

$$\lambda_k = (\tau_k - t) - \rho_{(k-R_\{0\})}\} \tag{17}$$

Here a pack counts as ready only when full ($\sigma \geq 0.995$), and ρ includes the constant-voltage taper above $\sigma = 0.9$ at half power. A pack whose matched demand has $\lambda_k \leq 15$ min is *urgent*; urgency is latched until the pack is full, and urgent packs charge at full power, PV first and grid for the remainder. Every other pack charges from PV only, healthiest first, and packs below 0.85 SoH are limited to 0.6 kW when not urgent. Greedy charging (GR), the baseline, charges every non-full pack at 1 kW as soon as it is docked.

D. Retirement Rule

Retirement depth is a design decision, not a fixed convention. Four rules are evaluated: fixed thresholds of 80%, 75% and 70% SoH, and a *demand-matched quota* (Q) that retires a pack below 0.70, or below 0.80 once the number of other packs under 0.80 reaches $Q = M_{\text{low}} + 2$, where M_{low} is the number of low-need riders. The quota ties the population of weakened packs to the demand that can use them. The sizing loop selects the rule together with N_s and P_{pv}; in what follows "DASS" denotes the health-aware dispatch with SA charging, and the retirement rule is stated explicitly.

E. Analytical Properties

Proposition 1 (wear-levelling maximises fleet fade). Let N packs share stress factor S and cycle fade $q_i = S\,A_i^z$ with $0 < z < 1$, and let the total throughput $\Sigma A_i = A$ be fixed. The fleet fade $F = S\Sigma A_i^z$ is Schur-concave in $(A_1, \ldots, A_N)$; it is maximised by the equal allocation, $F_{\max} = S\,N^{1-z}A^z$, and any allocation that majorises another has lower total fade.

Proof. $g(a) = a^z$ is concave for $0 < z < 1$, so $\Sigma g(A_i)$ is Schur-concave by Karamata's inequality [22]. The equal vector is majorised by every vector with the same sum, which gives the maximum. ∎

The result has a direct operational reading: a dispatcher that equalises cycling across packs (wear-levelling) cannot reduce, and generally increases, the total capacity the fleet loses. It may delay the *first* retirement, but it buys that by ageing everyone. Any benefit of health-aware dispatch must therefore come from matching capacity to need, not from redistributing fade. A Monte Carlo check with 1600 Dirichlet allocations of 46 packs confirmed the bound: the fade ratio $F/F_{\max}$ ranged from 0.710 to 1.0000, never exceeding one.

Proposition 2 (conditional optimality of assortative matching). Two riders with random needs N_1, N_2 receive two packs with deliverable energies $U_1 \geq U_2$. With $S_j(u) = \Pr(N_j > u)$, the assortative assignment ($U_1 \to 1$, $U_2 \to 2$) has fewer expected shortfalls than the swapped one iff

$$\Pr(U_2 < N_1 \leq U_1) \geq \Pr(U_2 < N_2 \leq U_1) \tag{18}$$

Proof. Expected shortfalls are $S_1(U_1) + S_2(U_2)$ and $S_1(U_2) + S_2(U_1)$; their difference is the difference of the two interval probabilities in (18). ∎ *Corollary:* if $U_2 \geq q_{1-\varepsilon}(N_2)$, the right side of (18) is at most ε, so assortative matching loses at most ε and gains the full mass of N_1 in $[U_2, U_1]$. When the weak pack does *not* cover the low-need rider's high quantile, matching can lose. This is exactly why DASS gives a weak pack to a low-need rider only when $U_i \geq \tilde{n}_r$, and why best-fit on the *mean* need (a weak pack covering only $\hat{n}_r$) fails in Section VII.

F. Algorithm and Complexity

TABLE II
DASS PROCEDURE PER 5-MIN STEP

Step	Operation
1	For each rider whose segment starts at t: add need; if $U < 1.1 \cdot$need, enqueue (planned).
2	Discharge riding packs; update T, q^{cyc}, q^{cal}; if $\sigma \leq \sigma_{res}$ mid-segment, enqueue (unplanned).
3	For each queued rider: compute $\tilde{n}_r$; select i^* by (15)–(16); else partial pack with $U \geq \tilde{n}_r$; else wait (≤ 20 min).
4	On deposit: update quota count; retire and replace if the retirement rule fires.
5	List demand instants τ_k (4-h look-ahead) + 2-pack buffer; sort non-full packs by time-to-full ρ; latch urgent packs with $\lambda \leq 15$ min (17).
6	Urgent packs: full power, PV then grid. Others: PV only, healthiest first, derate if SoH < 0.85.
7	Integrate charging energy, temperature (6) and ageing (4)–(5); log energy, swaps, waits, SoH.

With N packs and M riders, steps 1–2 cost $O(M)$, dispatch costs $O(N)$ per request, the retirement quota $O(N)$ per deposit, and step 5 is dominated by sorting, $O(N \log N + M \log M)$. A 5-year run of $T = 525{,}600$ steps therefore costs $O(T(N \log N + M \log M))$ time and $O(N + M + K)$ memory for K scheduled segments. The JIT-compiled kernel completes one five-year run in about 1.3 s on one CPU core, which is what makes the Monte Carlo sizing grid tractable; the policy itself is a few sorts per step and fits a kiosk controller.

VI. SIMULATION SETUP

The simulator was written in Python 3 (NumPy, SciPy, pandas) with the time-stepping kernel compiled by Numba and run on a two-core Linux workstation; figures use Matplotlib. Table III lists the parameters and labels each as sourced (S), assumed (A) or calibrated (C). Policies are named *dispatch–charging–retirement*: HB = health-blind FIFO, WL = wear-levelling (least-cycled full pack), HA = need-stratified assortative; GR = greedy, SA = solar-aware; 80/75/70 = fixed threshold, Q = quota. HB–GR–80 is the industry baseline. Two further controls isolate the dispatch design: MU applies best-fit to *every* rider on the *mean* need, and BF applies best-fit to every rider on the 95th-percentile need.

A worst-year constraint estimated from few replicates is optimistic for the design that happens to minimise cost, so the study uses three disjoint seed sets. (S1) *Screening*: 3,010 runs over $N_s \in \{12, \ldots, 24\}$, $P_{pv} \in \{0, 5, \ldots, 30\}$ kWp and 19 policies with five replicates (seeds 101–105). (S2) *Verification*: for each policy family, screened candidates are re-simulated in order of

cost with 20 new seeds (201–220) until the cheapest design that meets (13) on the 20-seed estimate is found and no cheaper candidate remains (960 runs). (S3) *Reporting*: all figures and tables of Section VII use a third set of 20 seeds (301–320) never used for selection, comprising the full factorial at the selected health-aware design (N_s = 18, P_{pv} = 10 kWp; 27 policies) and every family's selected design (640 runs). Each seed fixes the weather, rider draws and pack spread, so all policy and design differences are paired comparisons, reported as mean ± 95% confidence interval (CI) with paired *t*-tests.

TABLE III
SIMULATION PARAMETERS

Parameter	Value	Type
Time step / horizon	5 min / 5 years	A
Riders (RH / DL / PS)	12 / 10 / 8	A
Daily distance (RH / DL / PS)	140 / 100 / 30 km	S [1]
Energy intensity	0.032 kWh/km	S [1]
Pack	16S LFP, 40 Ah, 2.048 kWh	A
Slot power, charging eff.	1.0 kW, 0.94	A
Cycle-fade law	(4), Wang et al.	S [20]
Calendar fade	(5), 2%/yr at 25 °C, 50% SoC	C [21]
Capacity / rate spread	σ = 1% / log-σ = 0.15	A [8]
Annual GHI, PV tilt	4.87 kWh/m²/d, 20°	S/C [18]
Grid tariff	INR 5/kWh flat, no demand charge	S [15]
Value of time	INR 65/h	S [2]
Pack / slot / PV price	INR 13,500/kWh, INR 25,000, INR 45,000/kWp	A
Discount rate	10% real	A

VII. RESULTS AND DISCUSSION

Unless stated otherwise, numbers are means over the 20 reporting seeds (S3) with 95% CIs; u^* denotes the worst-year unplanned-swap rate per 1000 swaps and EAC the system EAC in thousands of INR per year (kINR/yr).

A. How Wrong Is the Identical-Battery Planning Model? (Q1)

The planning model selects N_s = 16 slots and 5 kWp of PV, where it predicts an EAC of 694.4 ± 4.9 kINR/yr and u^* = 11.0. Simulating the *same* design with ageing, and with the industry policy the planning model implicitly assumes (HB–GR–80), raises the EAC by 102.5 ± 3.4 kINR/yr (+14.8%, $p < 10^{-20}$) and u^* to 32.1, about 2.9 times the prediction (Fig. 3(a)). The shortfall has a single dominant cause: the planning model never replaces a pack, whereas the ageing fleet retires 21.9 of its 46 packs inside five years at the 80% threshold, because Bhubaneswar ambient temperatures plus charging heat roughly halve the cycle life relative to the 25 °C basis of a datasheet rating (Section IV-C). The rider-time cost also rises, since a fleet whose mean SoH sits near 0.87 even after replacements forces more mid-segment swaps. A planning model with identical, non-ageing packs is therefore biased in the direction that matters to an investor: it understates cost and overstates service.

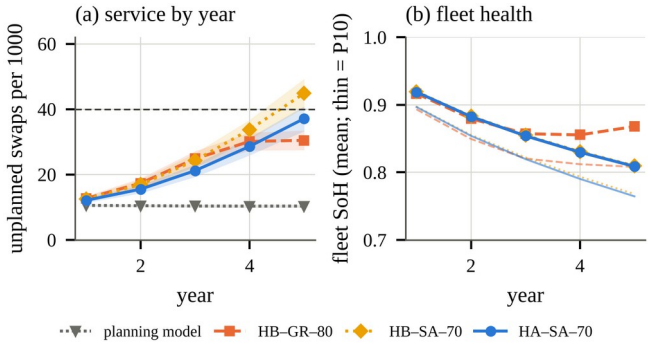


Fig. 3. Service and fleet health at N_s = 18, P_{pv} = 10 kWp (reporting seeds; bands are 95% CIs). (a) Unplanned swaps per 1000 by year; dashed line is the service limit. (b) Fleet mean SoH (markers) and 10th percentile (thin lines); the rise of HB–GR–80 after year 3 reflects replacements at 80% SoH.

B. Which Part of a Health-Aware Policy Creates Value? (Q2)

Table IV isolates the dispatch rule by holding charging and retirement fixed at the selected design. Health-aware dispatch lowers u^* in all eight charging–retirement combinations, by 5.8–10.4 per 1000 swaps (14–25% relative; every $p < 10^{-8}$). Its direct effect on cost is small and of either sign (−11.2 to +4.3 kINR/yr, i.e., within ±1.6% of EAC): the rule moves energy between riders, it does not create it. What changes is feasibility. Under health-blind FIFO dispatch, *every* retirement rule deeper than 80% violates the service limit at this design (u^* = 41.9–48.9); under health-aware dispatch every one of them satisfies it (u^* = 34.1–39.9). Fig. 3(a) shows where the difference arises: the two policies are indistinguishable in year 1, when all packs are healthy, and separate as the capacity distribution widens, which is exactly the regime in which the identity of the pack matters.

TABLE IV
PAIRED DISPATCH EFFECT AT THE SELECTED DESIGN

Charging–retirement	HB EAC	HB u^*	HA EAC	HA u^*	ΔEAC (HA−HB)	Δu^*	ΔEAC (WL−HB)
GR–80	810.3	30.5	810.6	23.0	+0.3	−7.5	+1.1
GR–75	740.8	44.9†	742.1	34.5	+1.3	−10.4	+1.2
GR–70	709.8	48.9†	708.1	39.9	−1.7	−9.0	+1.1
GR–Q	729.6	41.9†	718.4	35.6	−11.2	−5.8	+1.6
SA–80	767.9	30.8	772.2	24.0	+4.3	−6.4	+9.2
SA–75	706.7	42.6†	709.8	34.6	+3.1	−8.0	+7.4
SA–70	677.7	44.8†	680.2	37.1	+2.5	−7.7	+7.0
SA–Q	689.1	42.0†	689.2	34.1	+0.1	−7.8	+8.7

EAC in kINR/yr; u^* per 1000 swaps; † violates the limit of 40. N_s = 18, P_{pv} = 10 kWp, 20 reporting seeds. 95% CIs of ΔEAC are 0.5–4.0 kINR/yr; of Δu^*, 0.7–1.3.

The other two levers behave differently. Solar-aware charging is a pure energy lever: at fixed dispatch and retirement it lowers grid energy cost from 122.7 to 96.9 kINR/yr (−21%), raises the solar share of charging energy from 28% to 43%, and leaves u^* unchanged or slightly lower (Δu^* between −4.1 and +0.8). Retirement depth is the largest monetary lever: moving from 80% to 70% saves 90–92 kINR/yr, almost entirely through 15 fewer replacements, and costs 12.4–13.7 units of u^*. Health-aware dispatch is what lets an operator take that saving without breaching the service limit.

Negative results. Wear-levelling never helped: it raised EAC by 1.1–9.2 kINR/yr and u^* by up to 4.5 in all eight combinations (Table IV). This is the operational face of Proposition 1, which was confirmed numerically (Fig. 6(c): no allocation exceeded the equal-split fade). In the simulated fleet the fade penalty itself is small, because throughput dispersion across packs is low (Gini of equivalent full cycles ≤ 0.1 when no packs are replaced); at the 70% threshold, where replacements do not confound the comparison, wear-levelling increased total fleet fade by 0.3–0.6% ($p < 10^{-3}$). The naive best-fit control MU, which matches packs to the *mean* need of every rider, raised u^* by 12.3 and EAC by 11.8 kINR/yr against DASS with the same charging and retirement, as Proposition 2 predicts. Best-fit on the 95th-percentile need for every rider (BF) achieved a *lower* u^* (−3.7) than the stratified rule but at 9.2 kINR/yr more, because tight matching concentrated wear and triggered more quota retirements. Finally, the demand-matched quota did not beat a plain 70% threshold on cost: with health-aware dispatch it was 9.0 kINR/yr dearer, buying 3.0 units of service headroom. It is a conservative option, not a dominant one.

C. Joint Sizing Under the Service Constraint (Q3)

Fig. 4 shows the screening-stage cost and service curves with the verified optima, and Table V reports every selected design on the reporting seeds. The best health-aware design is HA–SA–70 with 18 slots and 10 kWp (EAC 680.2 ± 3.2); the best health-blind design that meets the same constraint is HB–SA–Q with 20 slots and 15 kWp (703.2 ± 5.5). The paired difference is −23.0 ± 3.6 kINR/yr (−3.3%, $p < 10^{-10}$). The composition of that gap is instructive: the health-aware design needs 2 fewer slots (−9.6 kINR/yr), 5 kWp less PV (−27.0) and fewer replacements (−17.5), but buys 28.1 kINR/yr more grid energy. The health-blind operator meets the constraint only by holding a larger, solar-heavier inventory and retiring packs earlier; the health-aware operator meets it by sending the right pack to the right rider. Note that the health-aware family's winner uses the fixed 70% threshold; DASS with the quota rule (16 slots, 10 kWp) is also feasible and 4.7 kINR/yr dearer.

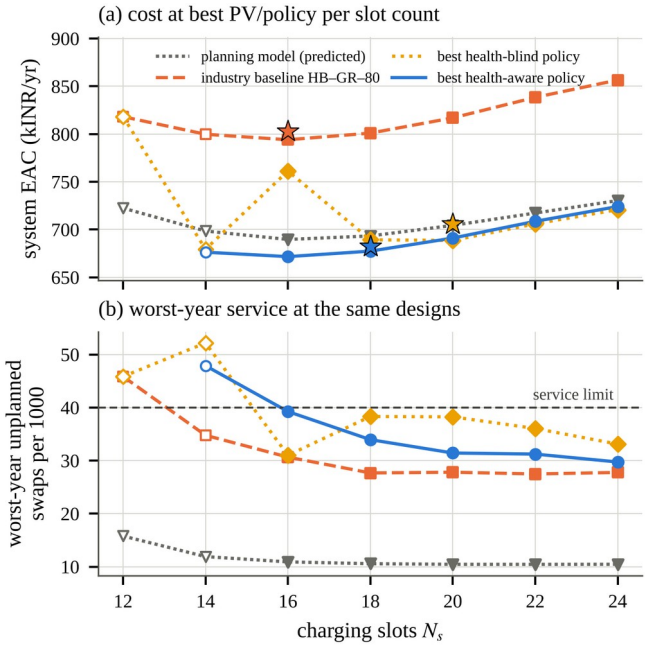


Fig. 4. Sizing curves from the screening stage (5 seeds): for each slot count, the cheapest PV size and policy within each family; filled markers meet the service constraints, hollow ones do not. Stars are the verified optima (20 seeds). (a) System EAC. (b) Worst-year unplanned swaps.

TABLE V
SELECTED DESIGNS ON THE REPORTING SEEDS

Design	Policy	N_s / P_{pv}	System EAC	Op. EAC	u^*	Solar	Repl.
Planning model (predicted)	PLAN	16 / 5	694.4 ± 4.9	532.8	11.0	16 %	0
Same design, realised (= industry baseline optimum)	HB–GR–80	16 / 5	797.0 ± 5.3	617.3	32.1	16 %	21.9
Best health-blind	HB–SA–Q	20 / 15	703.2 ± 5.5	538.5	38.8	58 %	5.2
DASS, quota	HA–SA–Q	16 / 10	684.9 ± 3.4	506.8	37.8	40 %	5.8
Best health-aware	HA–SA–70	18 / 10	680.2 ± 3.2	512.4	37.1	41 %	0.1

EAC in kINR/yr (system includes rider time). PV in kWp. Repl. = packs replaced in five years. The planning model selected the same design as the verified industry-baseline optimum; its first row is the model's own prediction, the second the simulated outcome with ageing.

Against the industry baseline the best health-aware design saves 116.8 ± 2.8 kINR/yr (14.7%). Honesty requires splitting this figure. The best health-blind design already captures 93.8 kINR/yr of it (80%) through solar-aware charging, deeper retirement and re-sizing, none of which needs per-pack health information. The increment attributable to health awareness is the remaining 23.0 kINR/yr. The baseline also carries the poorest kiosk availability of the selected designs (8.1 partial-pack hand-outs per 1000 swaps against 1.0), because it holds only 16 slots and charges greedily.

D. Operating Behaviour and Energy

Fig. 5 contrasts the two charging rules over two April days. Greedy charging tops up every returned pack immediately, so

evening returns draw the grid and the kiosk holds 15–18 full packs overnight, while midday PV is partly curtailed whenever the inventory is already full. Solar-aware charging defers non-urgent packs to the next solar window, keeps a smaller full inventory overnight, and still absorbs most of the available PV. Two side effects belong in any deployment discussion. First, just-in-time charging concentrates grid draw before the morning and evening demand waves, so the share of grid energy drawn between 18:00 and 24:00 is 40–43% under SA against 43% under greedy charging; under the flat Odisha EV tariff [15] this is costless, but it would not be under a demand charge. Second, deferral leaves packs at lower SoC for longer, which reduces calendar fade (−5% at the 70% threshold), while midday charging at higher ambient temperature raises cycle fade under the 80% rule (+7%); the net ageing effect of SA is small and of mixed sign.

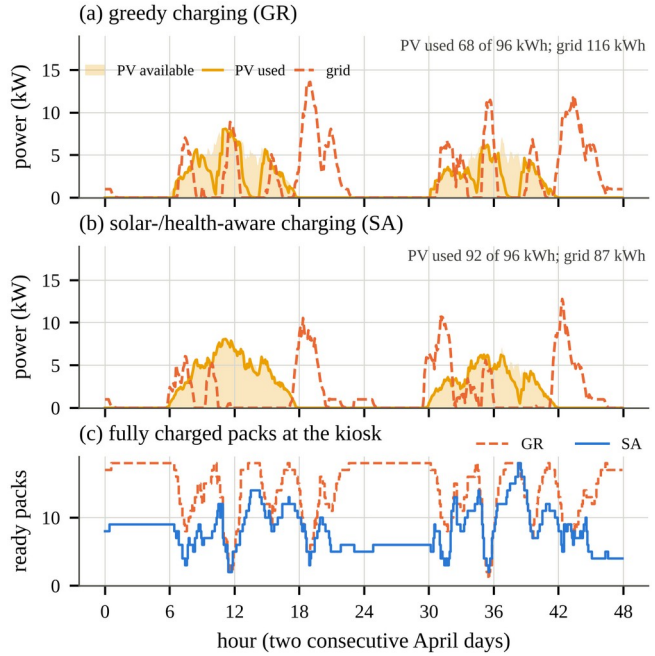


Fig. 5. Two-day operating trace (reporting seed 301, days 100–101) at 18 slots and 10 kWp. (a) Greedy charging. (b) Solar-/health-aware charging. (c) Fully charged packs at the kiosk.

E. Sensitivity and Robustness

Fig. 6(a) recomputes the paired saving of the selected health-aware design over the selected health-blind design under ten cost perturbations on the reporting seeds. The saving stays positive in every case, from 2.3% (value of time doubled, which rewards the health-blind design's larger inventory) to 5.4% (no residual value for packs), and every 95% CI excludes zero. It is largest when only operator cost counts (4.8%), because the health-aware design's advantage is capital, not rider time. Fig. 6(b) re-optimises both families on the screening data for different service limits, and this is the most important qualification in the paper. The value of health awareness is a *shadow price of the service constraint*: 9.3% at a limit of 30, 5.3% at 35, 2.5% at 40, and nil at 45 or above, where a health-blind operator can retire deeply without breaching the limit. At a limit of 25, no health-blind design in the search space is feasible, while BF at 20 slots and 20 kWp meets it at 728.6 kINR/yr (screening estimate). An operator that promises riders consistent range should therefore value per-pack health telemetry; one that does not need not.

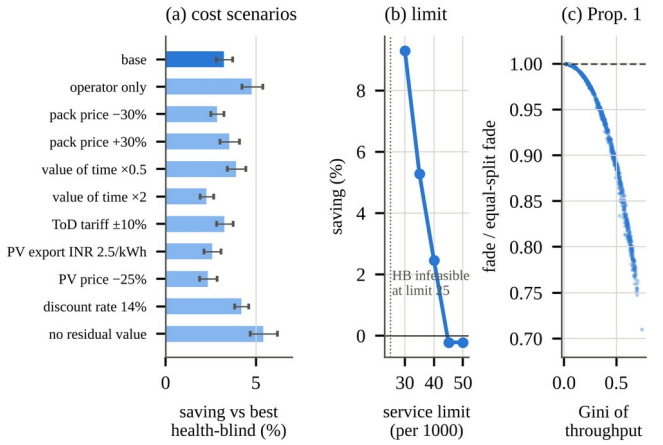


Fig. 6. (a) Paired EAC saving of the best health-aware over the best health-blind design under cost perturbations (reporting seeds, 95% CI). (b) Saving after re-optimising both families for different worst-year service limits (screening data). (c) Proposition 1: fleet fade relative to the equal-split maximum for 1600 random throughput allocations of 46 packs.

F. Corrections Made During the Study

Two errors were found during development and are reported because both would have inflated the result. (i) In the first version of the SA rule, a pack ceased to count as needing charge once it crossed the 95% dispatch threshold, so at night packs were dispensed at 95% rather than 100% SoC. This raised u^* by 7–14 per 1000 for every SA policy and made health-blind SA policies look worse than they are. The corrected rule (Section V-C) charges urgent packs to full and accounts for the constant-voltage taper. With the flawed rule the apparent advantage of the health-aware design over the best health-blind one was 7.9%; after correction it is 3.3%. (ii) Selecting designs on five replicates and then checking them on 20 fresh seeds showed that the screened optima of *both* families violated the worst-year constraint out of sample ($u^* \approx 43$–44 against a screened 38–39), a winner's-curse effect of minimising cost over many noisy constraint estimates. The three-stage protocol of Section VI was adopted in response. All 2,762 runs made before these corrections are excluded from the reported results and archived with the code.

VIII. Comparative Analysis

Table VI positions the method against the closest prior work on the criteria a station developer would weigh. Numerical comparison with the cited studies is not meaningful, because their sites, pack sizes, tariffs and objectives differ; the comparison is therefore methodological. The distinguishing features are the per-pack electro-thermal ageing inside the sizing loop, the rider-conditional dispatch, the worst-year service constraint and the paired three-stage Monte Carlo protocol. The price is computational: each design evaluation is a 5-year, 5-min simulation rather than a single LP solve, and the results inherit every assumption of the ageing and demand models.

TABLE VI
METHODOLOGICAL COMPARISON

Criterion	[3]	[5]	[14]	This work
Battery ageing	Not modelled	Linear in power	SPM surrogate	Arrhenius cycle + calendar, thermal
Pack heterogeneity	No	No	Yes	Yes (capacity, rate)
PV integration	Yes	No	No	Yes, calibrated
Station sizing	Yes	No	No	Slots, PV, retirement
Pack-to-rider dispatch	No	No	FIFO	Need-stratified
Uncertainty	Monte Carlo	Scenarios	Deterministic	MC, 3 seed sets
Service constraint	Not stated	Demand met	SoC at swap	Worst-year limits
Validation	Simulation	Simulation	Simulation	Simulation; inputs vs. atlas
Compute / design	LP-scale	MILP	20–40 s/step	≈1.3 s per 5-yr run

IX. KEY OUTCOMES AND LIMITATIONS

- **Planning bias.** An identical-battery model understated five-year EAC by 15% and the worst-year mid-segment swap rate by a factor of about 3 for the design it selected.
- **Service, not energy.** Health-aware dispatch reduced worst-year mid-segment swaps by 14–25% at near-zero direct cost and turned every deep retirement rule from infeasible to feasible.
- **Joint sizing.** Under a 40-per-1000 limit, the health-aware design met the constraint with 2 fewer slots and 5 kWp less PV, at 3.3% lower EAC; most of the saving over today's practice (80%) comes from solar-aware charging and retirement depth, which need no health data.
- **Practical significance.** The policy needs only BMS telemetry and subscription history that connected kiosks already collect, runs as a few sorts per 5-min step, and gives operators a quantitative case for advertising a range guarantee, the consumer-protection gap noted in [2]. Lower grid draw and fewer replaced packs also reduce emissions and recycling volume.

The limitations are those of a simulation study and are stated plainly. (a) No hardware validation is claimed; the ageing law is a published cell-level model [20] applied to packs with an assumed rate dispersion, and pack-level knee behaviour, cell imbalance and BMS SoH-estimation error (typically 2–5%) are not modelled. Estimation error would blur the capacity ranking on which dispatch relies and is the most likely way the benefit shrinks in practice. (b) Weather is synthetic, calibrated to long-run means; inter-annual extremes and heat waves are not represented beyond the AR(1) anomaly. (c) Rider behaviour is rule-based with three archetypes and a fixed segment structure; rider churn in response to service quality, which is the real economic penalty of a poor service record, is represented only through the service constraint. (d) A single kiosk is modelled; a network would let operators rebalance packs between kiosks, which may substitute for some of the effects shown here. (e) Costs are 2026 assumptions for India and are varied but not forecast; pack-price declines would reduce the value of deep retirement. (f) The sizing search is a grid with discrete steps; finer grids could shift optima by one step without changing the conclusions.

X. CONCLUSION

This paper replaced the identical-battery assumption in solar-integrated swap-station sizing with per-pack electro-thermal ageing and asked what that changes. Three answers emerged. First, the conventional model is not a harmless simplification: for an Indian two-wheeler kiosk it understates cost by about a seventh and mid-trip swaps by about a factor of three. Second, once packs differ, the question of which pack goes to which rider has a clear best answer, assortative matching on quantile need, whose value appears as service quality rather than energy cost; wear-levelling, the intuitive alternative, is provably counter-productive for total fleet fade under a concave fade law. Third, when sizing and scheduling are solved together under a worst-year service constraint, health awareness lets a smaller station retire packs later and still keep its promise to riders, saving 3.3% against the best health-blind design at a 40-per-1000 limit and more at tighter limits, while most of the saving over current practice comes from solar-aware charging and retirement depth. Future work will add BMS estimation error and pack-level knee ageing, extend the model to a network of kiosks with inter-station rebalancing, replace the rule-based dispatcher with a learned policy constrained by the propositions above, and validate the scheduler on an instrumented kiosk in Bhubaneswar.

ACKNOWLEDGMENT

The authors thank the Research and Development Division of Tekstein Solutions Private Limited for computing support. [Sponsor acknowledgments, if any, to be added by the authors.]

REFERENCES

[1] N. Dash and A. Bandivadekar, “Cost comparison of battery swapping, point charging, and ICE two-wheelers in India,” Int. Council Clean Transp. (ICCT), Washington, DC, USA, Working Paper 2021-38, Nov. 2021.

[2] P. Gode, S. Kohli, and J. Callahan, “Battery swapping for electric two-wheelers in India: Strategy hinterlands,” ICCT Working Paper 2022-17, with NITI Aayog, New Delhi, India, May 2022.

[3] N. Manoj, B. Vigneshkumar, A. G. Kumar, and M. R. Sindhu, “Optimal sizing and operation of PV based battery swapping station considering EV uncertainties for an Indian residential community,” in *Proc. 7th Int. Conf. Comput. Appl. Elect. Eng.—Recent Advances (CERA): Sustainable Transportation Systems*, 2023, IEEE Xplore doc. 10455128.

[4] B. Vigneshkumar and A. G. Kumar, “Optimal sizing and operation of solar PV powered EV battery swapping station for Indian petroleum retail outlet,” in *Proc. 3rd Int. Conf. Intell. Comput. Instrum. Control Technol. (ICICICT)*, 2022, doi: 10.1109/ICICICT54557.2022.9917711.

[5] M. R. Sarker, H. Pandžić, and M. A. Ortega-Vazquez, “Optimal operation and services scheduling for an electric vehicle battery swapping station,” *IEEE Trans. Power Syst.*, vol. 30, no. 2, pp. 901–910, Mar. 2015, doi: 10.1109/TPWRS.2014.2331560.

[6] Y. Zheng, Z. Y. Dong, Y. Xu, K. Meng, J. H. Zhao, and J. Qiu, “Electric vehicle battery charging/swap stations in distribution systems: Comparison study and optimal planning,” *IEEE Trans. Power Syst.*, vol. 29, no. 1, pp. 221–229, Jan. 2014.

[7] D. Renga, G. Centonze, and M. Meo, “Renewable powered battery swapping stations for sustainable urban mobility,” in *Proc. IEEE Int. Smart Cities Conf. (ISC2)*, 2022, pp. 1–7.

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

[9] K. A. Severson, P. M. Attia, N. Jin, N. Perkins, B. Jiang, Z. Yang, M. H. Chen, M. Aykol, P. K. Herring, D. Fraggedakis, M. Z. Bazant, S. J. Harris, W. C. Chueh, and R. D. Braatz, “Data-driven prediction of battery cycle life before capacity degradation,” *Nature Energy*, vol. 4, pp. 383–391, 2019.

[10] R. Rohde, A. Belea, M. Seibert, O. Gielen, A. M. Abdelalim, B. Ashok, and M. E. Farrag, “Developing role model of PV powered battery swapping stations for e-scooters in urban regions,” in *Proc. IEEE Int. Symp. Environment-Friendly Energies Appl. (EFEA)*, 2022, doi: 10.1109/EFEA56675.2022.10063829.

[11] X. Chen, Y. Yang, J. Wang, J. Song, and G. He, “Battery valuation and management for battery swapping station,” *Energy*, 2023 (preprint arXiv:2302.14291).

[12] “Battery swapping station management with heterogeneous battery functionality degradation,” *IEEE Trans. Smart Grid*, 2024–2025, IEEE Xplore doc. 10506804.

[13] A. A. Shalaby, H. Abdeltawab, and Y. A.-R. I. Mohamed, “Model-free dynamic operations management for EV battery swapping stations: A deep reinforcement learning approach,” *IEEE Trans. Intell. Transp. Syst.*, 2023, doi: 10.1109/TITS.2023.3264437.

[14] R. Li, Z. Chen, Z. Zhang, R. Guo, Z. Sun, J. Yao, and J. Ma, “Degradation-aware model predictive control for battery swapping stations under energy arbitrage,” arXiv:2510.07902, 2025.

[15] Odisha Electricity Regulatory Commission, Retail Supply Tariff orders for FY 2025–26 and FY 2026–27 (EV charging stations billed at a single-part INR 5/kWh; ToD tariff from 1 Jul. 2026), as reported by Mercom India, Apr. 2025 and 2026.

[16] J. A. Duffie and W. A. Beckman, *Solar Engineering of Thermal Processes*, 4th ed. Hoboken, NJ, USA: Wiley, 2013.

[17] D. G. Erbs, S. A. Klein, and J. A. Duffie, “Estimation of the diffuse radiation fraction for hourly, daily and monthly-average global radiation,” *Solar Energy*, vol. 28, no. 4, pp. 293–302, 1982.

[18] ESMAP/World Bank, “Global Solar Atlas,” site value for Bhubaneswar (annual GHI 1778 kWh/m^2, specific yield 1477 kWh/kWp) as compiled by Qbits Energy, “Solar irradiance data in India: How to compare states,” Oct. 2026.

[19] Japan Meteorological Agency, Tokyo Climate Center, “Monthly normals data: Bhubaneshwar, India (station 42971),” accessed Oct. 2026.

[20] J. Wang, P. Liu, J. Hicks-Garner, E. Sherman, S. Soukiazian, M. Verbrugge, H. Tataria, J. Musser, and P. Finamore, “Cycle-life model for graphite-LiFePO$_4$ cells,” *J. Power Sources*, vol. 196, no. 8, pp. 3942–3948, 2011.

[21] M. Naumann, M. Schimpe, P. Keil, H. C. Hesse, and A. Jossen, “Analysis and modeling of calendar aging of a commercial LiFePO$_4$/graphite cell,” *J. Energy Storage*, vol. 17, pp. 153–169, 2018, doi: 10.1016/j.est.2018.01.019.

[22] A. W. Marshall, I. Olkin, and B. C. Arnold, *Inequalities: Theory of Majorization and Its Applications*, 2nd ed. New York, NY, USA: Springer, 2011.