HPU2. Nat. Sci. Tech. Vol 05, issue 02 (2026), 41-55

HPU2 Journal of Sciences:
Natural Sciences and Technology

Journal homepage: https://sj.hpu2.edu.vn

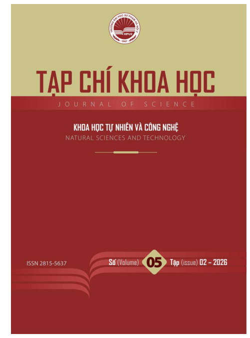

Article type: Research article

Multi-station smart charging in Vietnam: a quadratic-programming approach under the 2025 retail-tariff revision

Kim-Thuy Dinh Thi*

Hanoi Pedagogical University 2, Phu Tho, Vietnam

Abstract

We formulate day-ahead smart charging of an electric-vehicle fleet across multiple public charging stations in Vietnam as a convex quadratic program. The objective combines time-of-use tariff cost, quadratic peak-shaving, and ramp-smoothing; the constraints enforce per-EV connection-window bounds, per-EV energy delivery, per-station transformer caps, and feeder-level caps where stations share a primary distribution feeder. Station connector-mix and rate caps are drawn from OpenChargeMap and co-located PV from NASA POWER; arrival traces are transplanted from the real Shenzhen UrbanEV benchmark and recalibrated to the 2024–2025 VinFast fleet. The tariff reflects the 2025 retail-tariff restructuring (Decision 14/2025/QD-TTg). Across five seeds: F1, single-station Smart charging at α=0 cuts tariff cost 21.6% over uncontrolled charging, trading down to about 10% for a two-thirds peak-power reduction at higher peak-shaving weights; F2, on a 500-EV five-station feeder at the binding 1 MW cap the decentralised baselines are feasible in only 3 of 5 seeds while the coordinated solver stays feasible in all 5, at a 0.3–0.9% cost edge; F3, the 2025 flat-midday bracket cuts cost 14.7%; F4, OSQP solves the 500-EV instance in about 2 s. The study is a reproducible proof of concept for Vietnamese multi-station smart charging.

Keywords: quadratic programming, electric-vehicle charging, feeder coordination, time-of-use tariff, Vietnam

1. Introduction

Vietnam's electric-vehicle market expanded sharply through 2022–2025, led by VinFast's domestic vehicle production and the V-Green public charging network, which surpassed 150,000 charging ports by mid-2025. The Revised Power Development Plan VIII (Revised PDP8) [1], approved by Prime Minister's Decision 768/QD-TTg of 15 April 2025, targets a solar+wind share of 39.5 – 47.2 % of installed

* Corresponding author, E-mail: dinhthikimthuy@hpu2.edu.vn

https://doi.org/10.56764/hpu2.jos.2026.5.2.41-55

Received date: 25-5-2026 ; Revised date: 23-6-2026 ; Accepted date: 29-7-2026

generation capacity by 2030; the companion retail-tariff restructuring (Decision 14/2025/QD-TTg of 29 May 2025) [2] reshapes the time-of-use brackets that drive operator-side smart-charging decisions, which means EV charging will increasingly draw from a supply mix that is variable on hourly time-scales. At the same time, local distribution-transformer capacity in dense urban areas (Hanoi Old Quarter, HCMC District 1) is already a binding constraint on night-time charging at scale.

We address the following day-ahead question: given a real public-station inventory, a time-of-use tariff reflecting the 2025 retail-tariff restructuring, and station-level transformer caps, what is the smart-charging schedule that minimises tariff cost while limiting peak net-load excursions? We pose this as a compact convex quadratic program. The headline configuration uses a 24-hour horizon at hourly resolution; the solver-scaling experiments use horizons of up to 96 slots. Fleet sizes range from 50 to a few thousand vehicles distributed over a small number of public stations on a shared distribution feeder.

Global smart-charging literature. Quadratic-programming formulations of EV charging are well established. Sojoudi and Low [3] and Gan, Topcu and Low [4] are early references for convex EV-charging optimisation and decentralised protocols. Feeder-level and aggregator-level coordination under transformer-loading constraints has been studied by Knezović et al. [5] for hierarchical DSO/aggregator coordination across low-voltage networks, and more recently by van Huffelen et al. [6] for online scheduling of flexible EV charging under grid constraints. Multi-station and station-cluster coordination specifically has been treated by Liu et al. [7], who schedule across charging stations under a spatiotemporal demand distribution, and by Liao et al. [8] for a battery-swapping station cluster coupled to the grid; neither targets the Vietnamese tariff or public-station-inventory setting. Our QP and feeder-cap formulation is methodologically standard against this backdrop; the contribution we claim is positional, rather than algorithmic.

Table 1. Positioning against representative coordinated and multi-station EV-charging studies. “Feeder cap” denotes an explicit shared distribution-transformer / feeder limit coupling several stations; TOU = time-of-use tariff; “—” = not modelled.

Study	Optimisation method	Feeder/grid constraint	Tariff model	Scope & key finding
Knezović et al. 2021 [5]	Hierarchical DSO/aggregator optimisation	LV-network voltage/loading limits	Dynamic grid tariff	Multi-node LV coordination framework; non-Vietnamese
Liu et al. 2025 [7]	Collaborative multi-station scheduling	Cross-station power coordination	TOU	Multi-station spatiotemporal coordination; non-Vietnamese
Vu et al. 2024 [9]	MINLP (siting)	Network siting limits	—	Vietnamese station *placement*, not operation
Nguyen et al. 2025 [13]	LP + MPC, Bertsimas–Sim robust	Single-station cap	TOU + price uncertainty	Vietnamese-affiliated; single station on Caltech ACN data; 12–19% savings
This work	Convex QP (peak-shaving + ramp) with cap-duals	Shared feeder cap across stations	2025 Vietnamese retail TOU restructuring	Multi-station Vietnamese feeder; feasibility-driven coordination + tariff sensitivity; open artefact

Vietnamese EV optimisation literature. Five recent studies frame the Vietnamese context. Vu, Ha and Vu [9] solve a mixed-integer nonlinear *placement* problem for HCMC, expressly noting that operational studies are difficult in the Vietnamese context due to centralised grid control. A 2025 digital-twin study [10] optimises a single deployment on Hanoi university-campus using non-public sensor data. Minh et al. [11] provide a HOMER-based technoeconomic backdrop for PV-coupled charging in Hanoi, Da Nang, and HCMC, but stop short of scheduling. A separate single-station scheduling study with PV and storage integration [12] treats one station in isolation. Closest in scope to the present work is Nguyen, Pham and Do [13], a Vietnamese-affiliated (VinUniversity) single-station LP+MPC study under price uncertainty with a Bertsimas–Sim robustness layer, reporting sub-five-second solves and 12–19 % savings against a FCFS baseline on Caltech ACN data. We differ in three ways: (i) we coordinate across *multiple public stations sharing one distribution feeder* rather than scheduling a single station; (ii) our inputs are anchored on the Vietnamese *public*-station inventory exposed via OpenChargeMap [14] (whose VN coverage is sparse, motivating the modelled feeder topology of §3.2) rather than the Caltech ACN data; and (iii) the tariff schedule is the 2025 retail-tariff restructuring of Decision 14/2025/QD-TTg [2], which the Vietnamese references either predate or overlap with but do not explicitly study. A recent survey of quadratic programming and quadratically constrained quadratic programming [15] catalogues applications in this areas. Table 1 locates this work against the most comparable coordinated/multi-station and Vietnamese studies.

Contributions. Our contributions are positional: we apply a methodologically standard convex-QP layer to a Vietnamese setting that, to our knowledge, has not been treated reproducibly before.

1. A documented, runnable Vietnamese smart-charging pipeline using the OpenChargeMap [14] public-station inventory (with the disclosed caveat that VN coverage is sparse, $\sim$ 6 stations), the Shenzhen UrbanEV benchmark [16] arrivals transplanted via the rule in section 3.1, and NASA POWER irradiance for PV. The whole pipeline is exercised by the F1–F4 reproduction.
2. A multi-station formulation in which several public stations share one distribution-feeder cap, contrasting with the single-station scope of the cited Vietnamese references (10–13). The QP and feeder-cap constraint follow the global aggregator/coordination literature exemplified by [5,6]; our claim is the Vietnamese-context application, not the formulation.
3. A tariff-sensitivity analysis under the 2025 retail-tariff restructuring (Decision 14/2025/QD-TTg) [2], which the Vietnamese references do not explicitly study.

Two findings are substantive beyond the application framing: in the binding-cap regime, multi-station coordination is primarily a *feasibility* mechanism rather than a cost-reduction one (F2); and the Proposition 2 cap-duals expose *locational scarcity prices* that localise the binding load to the hour and feeder segment, which an operator can read as a transformer-reinforcement signal.

The remainder of the paper is organised in four parts. Section 2 (Experimental methods) defines the smart-charging quadratic program, proves convexity, and includes the Shenzhen-to-Vietnam transplant assumptions and the transformer-cap calibration. Section 3 (Results and discussion) reports the data pipeline, tariff savings, peak shaving, tariff sensitivity under the 2025 restructuring, and solver benchmarks at fleet sizes 50, 500, and 2000. Section 4 concludes and outlines stochastic and real-time extensions.

2. Model and experimental methods

This section presents the methodological core of the study: the convex quadratic program that schedules charging (§2.1–§2.7). The open data pipeline and case-study calibration that instantiate it for Vietnam are deferred to §3.1.

2.1. The operational scenario

We formalise the day-ahead decision faced by a city-scale charging operator at the close of business on the day before. The operator *commits* to a power profile for every electric vehicle it expects to host the next day, and that profile is the schedule delivered to each charger when the vehicle plugs in. The commitment is binding at the resolution of the planning slot $\Delta t \in \{15,30,60\}$ minutes, with power piecewise-constant within a slot. The decision rests on three exogenous inputs:

1. A *tariff schedule* τ_t in VND/kWh, published by EVN under the 2025 retail-tariff restructuring [2]. The operator pays τ_t for every kilowatt-hour drawn from the grid in slot t.
2. A *fleet manifest* $\{a_i, d_i, E_i, \bar{P}_i\}$ giving, for each EV i booked into a station the operator manages, the arrival slot a_i, the departure slot d_i, the energy delivery requirement E_i (kWh), and the per-vehicle rate cap $\bar{P}_i$ (kW) implied by the connector type.
3. A *network capacity envelope* $\bar{P}_k^{\text{tx}}$ (per-station transformer) and $\bar{P}_f^{\text{feeder}}$ (shared distribution-feeder), plus per-station baseline non-EV load $L_{k,t}$ and PV generation $S_{k,t}$.

The operator chooses, simultaneously for every EV and every slot, the kilowatt rate at which to charge. The choice is constrained by each EV's connection window, each EV's required energy delivery, the per-station transformer cap, and, where stations share a primary distribution feeder, a shared aggregate cap. The objective combines the operator's monetary cost (tariff energy bill), a quadratic penalty on the net load presented to the upstream grid (peak shaving), and a quadratic penalty on hour-to-hour rate changes per vehicle (ramp smoothing). The penalty weights α, β are operator-tuned knobs; their dimensions are chosen so that the three terms are on comparable monetary scales when $\alpha, \beta \sim O[1]$ VND kW^{-2}. The formulation below makes each of those choices explicit and the constraint set minimal.

2.2. Inputs and decision variable

Let $\mathcal{I}$ index EVs to be scheduled on the planning day, $\mathcal{K}$ index public charging stations, and $\mathcal{T} = \{0,1,\dots,T-1\}$ index discrete time slots of width Δt hours over a 24-hour horizon. The map $k: \mathcal{I} \to \mathcal{K}$ assigns each EV to a station and is taken as exogenous; we write $I_k = \{i: k(i) = k\}$.

Stations are partitioned into *feeders*: let $\mathcal{F}$ index the set of distribution feeders and let $f: \mathcal{K} \to \mathcal{F}$ map each station to its feeder, with $K_f = \{k: f(k) = f\}$ the stations on feeder f. Several public stations on the same primary distribution circuit share a single transformer; we capture this with a feeder-level cap $\bar{P}_f^{\text{feeder}}$. When every feeder contains exactly one station, the per-station cap $\bar{P}_k^{\text{tx}}$ and the feeder cap coincide. The decision variable is the per-EV power profile $x_{i,t} \geq 0$ in kilowatts. We write $\mathbf{1}[\cdot]$ for the indicator function.

Definition 1 (Inputs). *A smart-charging instance is the tuple* $(\mathcal{I}, \mathcal{K}, \mathcal{F}, f, T, \Delta t, \{a_i, d_i, E_i, \bar{P}_i\}_i, \{\bar{P}_k^{\text{tx}}, L_{k,\cdot}, S_{k,\cdot}\}_k, \{\bar{P}_f^{\text{feeder}}\}_f, \tau_{\cdot}, \alpha, \beta)$, *with* $a_i < d_i \leq T$, $E_i \geq 0$, $\bar{P}_i > 0$, $\bar{P}_k^{\text{tx}} > 0$, $\bar{P}_f^{\text{feeder}} > 0$, $\alpha, \beta \geq 0$, *and the per-EV* reachability condition $E_i \leq \bar{P}_i(d_i - a_i)\Delta t$, which is necessary for the single-EV subproblem to be feasible. Joint feasibility additionally requires the per-station and feeder caps to admit the aggregate energy demand in each slot, a condition we verify numerically on each instance.

Modelling choices. The EV-to-station map k is exogenous: we do not optimise over *which* station an EV uses. That is the upstream station-placement and routing problem (Vu et al. [9]), which we treat as solved. The slot resolution Δt is constant across the horizon; the formulation accommodates any constant Δt but does not handle non-uniform discretisation. The energy-delivery constraint is an equality rather than an inequality, matching the contractual setting where the operator agreed to deliver E_i kWh; allowing under-delivery (e.g. for top-up sessions) would replace it with $\sum_t x_{i,t}\, \Delta t \geq E_i^{\min}$ and forfeit the closed-form structure used in Lemma 1 below.

2.3. Problem (P)

Problem 1 (Smart-charging QP). *Given an instance per Definition 1, find* $x \geq 0$ *that*

$$\boxed{\min_{x\geq 0}\ \underbrace{\textstyle\sum_t \tau_t \left(\sum_i x_{i,t}\right)\Delta t}_{\text{tariff cost}} \ +\ \underbrace{\alpha \textstyle\sum_k \sum_t \left(\sum_{i\in I_k} x_{i,t} + L_{k,t} - S_{k,t}\right)^2}_{\text{peak shaving}} \ +\ \underbrace{\beta \textstyle\sum_i \sum_{t\geq 1}\left(x_{i,t} - x_{i,t-1}\right)^2}_{\text{ramp smoothing}}} \quad (1)$$

$$\text{subject to} \quad \begin{aligned} 0 \leq x_{i,t} &\leq \bar{P}_i \cdot \mathbf{1}[a_i \leq t < d_i] \quad \forall i,t, && (2)\\ \textstyle\sum_t x_{i,t}\ \Delta t &= E_i \quad \forall i, && (3)\\ \textstyle\sum_{i\in I_k} x_{i,t} + L_{k,t} &\leq \bar{P}_k^{\mathrm{tx}} \quad \forall k,t, && (4)\\ \textstyle\sum_{k\in K_f}\left(\sum_{i\in I_k} x_{i,t} + L_{k,t}\right) &\leq \bar{P}_f^{\mathrm{feeder}} \quad \forall f,t. && (5)\end{aligned}$$

The per-station (4) and feeder (5) caps are bound *through-power* drawn from the grid, so PV generation $S_{k,t}$ does not appear on the left-hand side; PV enters the formulation only through the peak-shaving term, which targets net load. If a deployment exports PV upstream (V2G), (4) and (5) would need a sign-relaxed extension; we do not treat that here (see §4).

Objective. The three terms in (1) play distinct roles.

Tariff cost is linear in x and directly accounts for the energy bill the operator owes to EVN at the close of the planning day. The factor Δt converts the per-slot power $x_{i,t}$ to the slot's energy $x_{i,t}\Delta t$, against which τ_t is priced. With $\alpha = \beta = 0$ the problem becomes a pure linear program and concentrates each EV's energy delivery into the cheapest reachable slots, subject to the per-EV rate cap and the per-station and feeder caps. That LP has all the structure of a multi-knapsack and provides the lower bound PerfectLP we use in F1.

Peak shaving is the primary reason for choosing a QP over an LP. The quadratic $\sum_{k,t}\left(\sum_{i\in I_k} x_{i,t} + L_{k,t} - S_{k,t}\right)^2$ penalises deviations of each station's net load from zero; minimising it tends to flatten station-level demand. Without it, the LP-driven schedule concentrates EV charging into the cheapest off-peak hours, producing the secondary peak that distribution operators aim to avoid. The weight α converts squared kilowatts into VND; its calibration is operator-discretionary and is the F1 axis we sweep.

Ramp smoothing discourages hour-to-hour swings in any one EV's rate. From the operator's view it is a comfort/equipment-wear proxy; from the QP view it is a regulariser that ensures the ramp-smoothing Hessian $\beta D^\top D$ has the same per-EV tridiagonal structure for every vehicle. With $\beta = 0$ the term vanishes and the Hessian reduces to the peak-shaving block $\alpha A^\top A$, which is block-diagonal in slot index t and couples EVs at the same station within each slot through the rank-one Gram outer product $\mathbf{1}_{I_k}\mathbf{1}_{I_k}^\top$. The F1 and F2 instances reported below use $\beta = 0$.

Constraints. The four constraint families have distinct roles.

(2) is the *connection window*: power can only be drawn while the EV is physically plugged in. With it, the schedule must fit the energy demand into whatever slots are available, which is what makes day-ahead smart charging worth solving rather than buying spot energy.

(3) is the *energy delivery* commitment, an equality binding the operator to honour the booking.

(4) is the *per-station transformer cap*: the aggregate charging power at station k in slot t, plus the station's non-EV baseline load $L_{k,t}$, may not exceed the transformer's rated through-power. The cap holds in every slot; typically, only a handful of slots bind.

(5) is the *shared feeder cap*. Several public stations on the same primary distribution feeder draw through one upstream transformer; the feeder cap is binding only on the *aggregate* load at the

feeder level. The relationship to (4) is hierarchical: a feasible schedule must respect both the per-station caps and the feeder cap; either can bind first depending on calibration.

When the feeder cap is binding, the joint optimum shifts load across stations to honour it; this cross-station shift is the mechanism that independent per-station solutions cannot reproduce and is the source of the F2 coordination gap.

Hard constraints versus soft penalties. The connection-window bounds (2), the energy-delivery equality (3), and the per-station (4) and feeder (5) caps are *hard constraints*: they carve out the feasible polytope and admit no violation, reflecting either physical impossibility (drawing power through an unplugged connector, or exceeding a transformer's rated through-power) or contractual obligation (the booked energy E_i). Peak shaving and ramp smoothing instead enter the *objective* as soft quadratic penalties weighted by α and β; they encode operator preferences, a smooth net load, and a steady per-EV rate, that are traded off against tariff cost rather than enforced.

2.4. Stacked-vector representation

Two linear operators recur throughout the analysis. Stack the decision variable in row-major order as $x \in \mathbb{R}_{\geq 0}^{|\mathcal{I}|T}$, indexed by (i,t). The *EV-to-station-slot summing matrix* $A \in \mathbb{R}^{|\mathcal{K}|T \times |\mathcal{I}|T}$ is defined by

$$A_{(k,t),(i,t')} \;=\; \mathbf{1}[k(i)=k]\cdot\mathbf{1}[t=t'],$$

together with the offset $b \in \mathbb{R}^{|\mathcal{K}|T}$, $b_{(k,t)} = S_{k,t} - L_{k,t}$. By construction $(Ax-b)_{(k,t)} = \sum_{i \in I_k} x_{i,t} + L_{k,t} - S_{k,t}$, so the peak-shaving term equals $\alpha \;\| Ax-b \|_2^2$.

The per-EV first-difference operator $D \in \mathbb{R}^{|\mathcal{I}|(T-1)\times|\mathcal{I}|T}$ is defined by

$$D_{(i,t),(i',t')} \;=\; \mathbf{1}[i=i']\,(\mathbf{1}[t'=t]-\mathbf{1}[t'=t-1]), \quad 1 \leq t \leq T-1,$$

so $(Dx)_{(i,t)} = x_{i,t} - x_{i,t-1}$ and the ramp-smoothing term equals $\beta \;\| Dx \|_2^2$. We use A, b, D in the convexity proof, the sparsity remark and the Schur-complement direction sketched in Section 4.

2.5. Convexity

Proposition 1 (Convexity of (P)). The objective in (1) is convex quadratic with positive semidefinite Hessian, and the feasible region defined by (2)–(5) is a polyhedron. Hence (P) is a convex quadratic program.

Proof. Use the operators A, b, and D from Section 2.4. The tariff term is linear in x. Expanding the peak-shaving term,

$$\alpha \;\| Ax-b \|_2^2 = \;\alpha\, x^\top A^\top A\, x \;-\; 2\alpha\, b^\top A\, x \;+\; \alpha\, b^\top b,$$

gives a quadratic in x with Hessian $2\alpha\, A^\top A$ and a constant offset $\alpha\, b^\top b$ that does not affect the minimiser; $A^\top A \succcurlyeq 0$ as a Gram matrix, and the linear term $-2\alpha\, b^\top A\, x$ adds to the QP's linear coefficient, giving the effective $\tilde{c} = c - 2\alpha\, A^\top b$ alongside the tariff coefficient. The ramp-smoothing term $\beta \;\| Dx \|_2^2$ has Hessian $2\beta D^\top D \succcurlyeq 0$ and zero offset.

The objective is therefore $x^\top H x + c^\top x + \text{const}$ with $H = \alpha A^\top A + \beta D^\top D \succcurlyeq 0$ as a sum of two PSD matrices, and is convex. Constraint (2) is a per-slot box; (3) is a linear equality; (4) and (5) are linear inequalities. The feasible set is therefore a polyhedron, and (P) is a convex QP. $\square$

Remark 1 (Sparsity). $A^\top A$ is block-diagonal in t with each $|\mathcal{I}| \times |\mathcal{I}|$ block formed from station-membership outer products $\oplus_{k \in \mathcal{K}} \mathbf{1}_{I_k} \mathbf{1}_{I_k}^\top$, so it has at most $T \sum_k |\, I_k|^2$ non-zeros (one $\mathbf{1}_{I_k}\mathbf{1}_{I_k}^\top$ block per slot). $D^\top D$ is block-diagonal in i with each $T \times T$ block tridiagonal (diagonal $\{1,2,\ldots,2,1\}$, sub/superdiagonal -1). The Hessian $H = \alpha A^\top A + \beta D^\top D$ inherits both sparsity patterns; this is the structure that the operator-splitting solvers of Section 2.8 exploit.

Remark 2 (Strict convexity is generic but not universal). $H = \alpha A^\top A + \beta D^\top D$ is only PSD on $\mathbb{R}^{|I|T}$, but strict convexity holds on the feasible affine subspace $\{v: \sum_t v_{i,t}\, \Delta t = 0\ \forall i\}$ pinned by the energy equality (3), so the minimiser of (1) is unique on the feasible set. At $\beta = 0$, a flat optimum can occur on a low-dimensional face; ties are resolved by OSQP's polish step.

2.6. KKT structure

When the transformer and feeder caps are slack at some feasible point, Slater's condition holds and the KKT conditions characterise the optimum. The F1 and F4 instances satisfy this assumption by construction; F2 reports both binding and slack regimes and we treat the dual quantities in Proposition 2 under the standard regularity assumption (linear independence of the active inequality gradients at the optimum).

Lemma 1 (KKT on the 1-EV interior). Consider a single EV alone at its station with $L_{k(i),t} = S_{k(i),t} = 0$ for all t, $d_i - a_i = T$, non-binding transformer and feeder caps, and $\beta = 0$, $\alpha > 0$. Assume the per-slot energy budget is large enough relative to the tariff spread, $\frac{E_i}{T\Delta t^2} \geq \frac{\max_t \tau_t - \text{mea}\ \ (\tau)}{2\alpha}$, so that the closed form below stays non-negative in every slot; assume additionally that the per-EV power cap admits the peak, $\bar{P}_i \geq \frac{E_i}{T\Delta t} + \frac{\left(\text{mean}(\tau) - \min_t \tau_t\right)\Delta t}{2\alpha}$, so that the closed form below stays at or below $\bar{P}_i$. Let $\mu \in \mathbb{R}$ be the multiplier on the energy equality (3). At an interior optimum, $x_{i,t} = \frac{(\mu - \tau_t)\,\Delta t}{2\alpha}$ with $\mu = \frac{1}{T} \sum_t \tau_t + \frac{2\alpha E_i}{T\Delta t^2}$. Without (8), the optimal pins against $x_{i,t} = 0$ in the spikiest slots and the closed form no longer applies.

Sketch. With $L = S = 0$ and only EV i at its station, the peak-shaving term reduces to $\alpha \sum_t x_{i,t}^2$. The Lagrangian on (3) (multiplier μ) gives stationarity $\tau_t \Delta t + 2\alpha\, x_{i,t} - \mu\, \Delta t = 0$ in each slot, so $x_{i,t} = (\mu - \tau_t)\Delta t/(2\alpha)$. Substituting into (3), $\sum_t x_{i,t}\, \Delta t = \Delta t^2/2\alpha\, (T\mu - \sum_t \tau_t) = E_i$, yielding $\mu = \text{mean}(\tau) + 2\alpha E_i/(T\Delta t^2)$. ☐

2.7. Algorithmic choice

Problem (P) has $|\mathcal{I}|\, T$ variables and $O(|\mathcal{I}|\, T + |\mathcal{K}|\, T)$ linear inequalities; the Hessian $H = \alpha A^\top A + \beta D^\top D$ inherits the block-diagonal-in-t sparsity of $A^\top A$ (per Remark 1) and the per-EV tridiagonal sparsity of $D^\top D$. Three solver families are available off-the-shelf for this regime, each with a different operating point in the speed/accuracy plane.

OSQP (operator-splitting, our default). OSQP [18] casts the QP in standard form $\min 1/2\, x^\top H x + c^\top x$ subject to $l \leq Bx \leq u$ and runs an alternating direction method of multipliers (ADMM) [19] iteration that alternates between (i) a linear system solve on the KKT matrix $\begin{bmatrix} H + \sigma I & B^\top \\ B & -\rho^{-1} I \end{bmatrix}$ and (ii) a closed-form projection onto the box $[l, u]$. The linear system has the same sparsity pattern at every iteration, so its Cholesky factorisation is computed once and reused; this is the source of OSQP's scaling on instances with $|\mathcal{I}|\, T \sim 10^5$. Two termination criteria are checked each iteration: a primal residual $\| Bx - z \|_\infty \leq \varepsilon_{\text{abs}} + \varepsilon_{\text{rel}} \ \| Bx \|_\infty$ and a dual residual of the same form on $Hx + c + B^\top y$. The optional *polish* step, run once after termination, identifies the active set and solves the equality-constrained QP on it directly to bring the working residuals close to machine precision; we enable polish in the experimental sweeps of Section 3.

All three are accessible through CVXPY's solver interface; tolerance settings (ε_{abs}, ε_{rel}, max_iter, polish) are passed through transparently. The F4 benchmark in Section 3.4 reports OSQP wall-clock at $\varepsilon_{\text{abs}} = \varepsilon_{\text{rel}} = 10^{-6}$ with polish enabled on the reference machine.

3. Experiments

3.1. Data and case study

All inputs are open enough to release a fully reproducible artefact. This section summarises the five data layers and the assumptions that bridge them to the formulation (1).

Public station inventory. OpenChargeMap [14] returns only ~ 6 Vietnamese public records under countrycode=VN as of mid-2025 (a 2026-05 snapshot included in the artefact); V-Green's $\sim 150{,}000$ ports are not pushed to OCM. We use the OCM snapshot for what it does support—real connector titles, observed PowerKW and LevelID/CurrentTypeID ratings, and the resulting per-vehicle rate-cap distribution—by mapping each connector to its nominal rating (AC 7/11/22 kW; DC 50/100/150 kW; CCS/CHAdeMO forced DC; OCM's PowerKW snapped to the nearest rung within the family). Because OCM coverage is too sparse to source a real multi-station feeder cluster, the five-station HCMC District 1 topology used in experiment F2 is an explicitly modelled scenario informed by the real connector mix and a representative port-count vector; coordinates and station identities are illustrative and reported as such.

Arrival–departure traces. No publicly available Vietnamese dataset provides per-vehicle arrival, departure, or energy demand at the hourly resolution that our formulation requires. We transplant from the Shenzhen UrbanEV benchmark [16], which records hourly occupancy across ~ 74 stations over 2022–2023; the dataset itself is released under CC0 1.0 (the companion Scientific Data article is CC BY-NC-ND 4.0). We derive the marginal arrival-rate density $\hat{p}(t)$ from hour-to-hour positive occupancy increments aggregated across all UrbanEV stations and days, and a pooled session-length distribution from the duration product; the F1–F3 fleets are sampled from these real distributions via the rule of the next paragraph.

Trace synthesis rule (UrbanEV-based). The pipeline uses UrbanEV's per-station hourly *arrival counts* $n_{k,t}^{\text{SZ}}$ (not occupancy, which folds in session length). Aggregating across Shenzhen stations and days and normalising, $\hat{p}(t) = \sum_{k,d} n_{k,t,d}^{\text{SZ}} \;/\; \sum_{k,d,t'} n_{k,t',d}^{\text{SZ}}$, is the empirical marginal arrival-rate density over the 24-hour day. The joint distribution of (arrival slot,session length) is captured by the empirical bivariate histogram $\hat{q}(t, \Delta)$ from completed sessions, with session length sampled conditional on the arrival slot. The F1–F3 fleets in this submission are sampled directly from these empirical UrbanEV densities $\hat{p}$ and $\hat{q}$; the arrival-shift sweep of F3 (§3.3) quantifies how sensitive the headline numbers are to the Shenzhen-to-Vietnam timing transfer.

The Vietnamese fleet size $|\mathcal{I}|$ is an exogenous experimental parameter (50 for F1, 500 for F2–F3, 50/500/2000 for F4). For each EV i we sample $a_i \sim \hat{p}$, $d_i - a_i \sim \hat{q}(\cdot|\, a_i)$, and assign i to a station with probability proportional to that station's OpenChargeMap port count. Per-EV energy demand E_i is drawn from a three-component mixture calibrated to the 2024–2025 VinFast model mix: probability 0.25 component $\mathcal{N}_{[4,18]}(12, 3^2)$ for VF 3, probability 0.60 component $\mathcal{N}_{[8,40]}(25, 6^2)$ for VF e34, and probability 0.15 component $\mathcal{N}_{[15,85]}(50, 12^2)$ for VF 8 (in kWh). Reachability $E_i \le \bar{P}_i(d_i - a_i)\Delta t$ is enforced by rejection-resampling.

Applicability to the Vietnamese setting. The Shenzhen marginal is a defensible stand-in because the schedule responds only to the coarse diurnal shape — daytime build-up, evening push, overnight trough — which is generic to dense-urban charging, not Shenzhen-specific. The behaviorally sensitive inputs are handled directly: arrival *timing* is stress-tested by the F3 shift sweep (Section 3.3, $\Delta a \in \{-2,\ldots,+2\}$ h), and per-session energy is recalibrated to the 2024–2025 VinFast mix. Validation against released Vietnamese traces remains future work (Section 4).

Time-of-use tariff. The tariff τ_t is encoded from EVN's published time-of-use schedule (peak/shoulder/off-peak brackets for industrial and commercial customers; The 2025 retail-tariff restructuring follows the schedule in EVN's price-decision instrument **Decision 1279/QD-BCT of 9 May 2025** (the legal vehicle implementing the structural changes initiated by Decision 14/2025/QD-TTg [2]);. Fig. 1 shows the resulting hourly schedule.

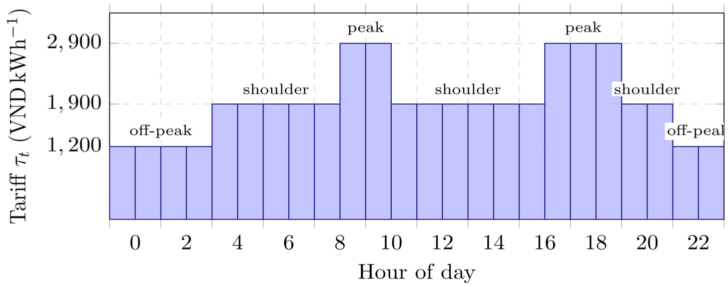

Figure 3. EVN time-of-use tariff brackets used in the experiments: off-peak 1,200 VND/kWh (22:00–04:00), shoulder 1,900 VND/kWh (04:00–09:00, 11:00–17:00, 20:00–22:00), and peak 2,900 VND/kWh (09:00–11:00, 17:00–20:00). Mean tariff 1,933 VND/kWh; peak-to-off-peak ratio $ ≈ 2.4×$. The midday-shoulder and evening-peak bands are the two intervals that the 2025 retail-tariff restructuring adjusts in scenarios 2025-flat-midday and 2025-steep-evening. Bracket boundaries are rounded to the nearest hour for the $\Delta t = 1$ h grid used in the experiments; the original EVN schedule has a 30-minute morning-peak start at 09:30.

PV generation and baseline load. Per-station PV output $S_{k,t}$ is computed from NASA POWER hourly point-irradiance at station coordinates, scaled by on-site nameplate capacity. For street-level urban hubs the nameplate is typically zero; PV enters meaningfully only for the small number of hubs with co-located rooftop installations. Per-station baseline load $L_{k,t}$ is synthesised from EVN quarterly load bulletins (Hanoi 2024 peak ≈3.2 GW; HCMC 2024 peak ≈5.8 GW), scaled to a per-feeder size of 1–5 MW typical of urban distribution transformers serving multi-station charging clusters. Because the per-feeder design standard is not consistently public, we treat $L_{k,t}$ as a parameter swept in the robustness analysis.

Transformer- and feeder-cap calibration. Two distinct caps enter (P): a per-station transformer cap $\bar{P}_k^{tx}$ and a shared feeder cap $\bar{P}_f^{feeder}$ aggregating across all stations in K_f. Per-station caps are calibrated based on the station's reported port count and connector mix, using the smaller of 300 kW and $1.2 \times \sum_{p\in ports(k)} \bar{P}_p$ as an upper bound for a single-station transformer. The shared feeder cap is the parameter that drives the multi-station behaviour: in the F2 sweep we take $\bar{P}_f^{feeder} \in \{1,2,5\}$ MW for a five-station feeder in HCMC District 1. Per-station caps are chosen large enough that the feeder cap binds first in the headline F2 cases. Empirical defence of either calibration would require EVN's per-feeder design records, which are out of scope; we therefore treat both as parameters of the sensitivity sweep rather than as fixed inputs.

3.2. Results and discussion

We evaluate (P) along four axes: a single-station baseline (F1), multi-station feeder coordination (F2), tariff sensitivity under the 2025 retail-tariff restructuring (F3), and a solver benchmark (F4). Each experiment is reproducible from a fixed random seed and the data pipeline of Section 3.1; the seeds and instance parameters appear with each result block.

3.2.1. F1, a single-station baseline

Setting. A 50-EV single-station instance over $T = 24$ hourly slots with a non-binding transformer cap $\bar{P}^{\text{tx}} = 1\text{MW}$, with arrivals, durations, and the three-component VinFast energy mixture drawn by the trace-synthesis rule of §3.1. We report 5-seed medians [min, max] for Smart(α), problem (P) at $\beta = 0$, swept over $\alpha \in \{0,1,10,10^2,10^3\}$VND kW^{-2} slot^{-1} — against uncontrolled ASAP (charge at $\bar{P}_i$from arrival until $E_i$is met), constant-rate Avg, and a window-/penalty-free tariff LP (PerfectLP) that lower-bounds achievable tariff cost. We report (i) total day-ahead tariff cost, (ii) station-level peak power $\max_t \sum_i x_{i,t} + L_t$, (iii) ramp $\max_t |\sum_i x_{i,t+1} - \sum_i x_{i,t}|$, (iv) per-EV charge completion at departure (should be exact).

F1 result (5 seeds, $n = 50$, $T = 24$, EVN nominal tariff). Smart sweep, median [min, max] tariff savings over ASAP: $\alpha = 0$ **21.6%** [20.4, 26.1], $\alpha = 1$ **21.6%** [20.4, 26.1], $\alpha = 10$ **10.3%** [3.9, 12.2], $\alpha = 100$ **2.9%** [−1.4, 4.2], $\alpha = 1000$ **2.0%** [−2.2, 3.5]. Peak power drops monotonically as α grows (representative seed: $\sim 298 \rightarrow 109\,\text{kW}$ from $\alpha = 0$ to $\alpha = 1000$). The $\alpha = 0$ case is the LP limit (tariff arbitrage only); the *Smart QP* regime is $\alpha \gtrsim 10$, where peak-shaving trades ~ 10 percentage-points of tariff saving for a roughly two-thirds reduction in peak power. At $\alpha = 1$ the schedule remains on the tariff-optimal face of the feasible polyhedron, hence the agreement with the $\alpha = 0$ savings. Avg delivers only $\sim 4\%$ over ASAP; PerfectLP bounds tariff-only savings at $\sim 36\%$. A $\beta > 0$ cell at $\alpha = 10, \beta = 1$ exercises ramp smoothing: savings $\sim 10.3\%$, peak $\sim 135\,\text{kW}$, maximum hour-to-hour ramp $\sim 50\,\text{kW}$. Median OSQP solve time per cell: 40–100 ms.

3.2.2. F2, a multi-station feeder coordination

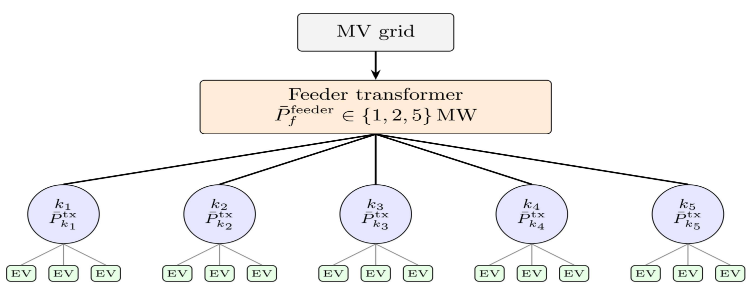

Figure 2. F2 topology: five public stations $k_1, \ldots, k_5$ share one primary distribution-feeder transformer with cap P̄ffeeder. Each station carries its own (larger) per-station cap P̄ktx which is non-binding by design in the F2 sweep so that the headline coordination claim is driven by the feeder constraint (5).

Setting. A modelled five-station feeder topology in HCMC District 1 (Figure 2), sparse OCM coverage prevents sourcing a real cluster, with a single aggregate cap $\bar{P}_f^{\text{feeder}} \in \{1,2,5\}$ MW. 500 EVs are distributed across stations by a representative port-count weight vector {6,4,8,5,4}; arrivals come from the real UrbanEV-derived sampler.

Comparisons. Three decentralised baselines span the range of plausible alternatives; the coordination claim is supported only if CoOpt improves on the best feasible one.

- IsoProp/ IsoDem: independent per-station QPs with the feeder cap split proportionally to port counts (IsoProp) or proportionally to per-station energy demand $\sum_{i \in I_k} E_i$ (IsoDem).
- IsoDual: a single-pass dual reallocation. We solve the IsoProp configuration, extract the per-station transformer cap shadow prices (Proposition 2) exposed by the solver, and reallocate the feeder cap by blending the dual signal with the proportional baseline ($w^{\text{dual}} = 1/2\, \lambda / \sum \lambda + 1/2\, w^{\text{prop}}$) so no station starves; then we independently solve with the reallocated caps.
- CoOpt: joint solution of (P) with (5) imposed.

We report (i) tariff-cost gap between each Iso* baseline and CoOpt (relative percentage), (ii) the maximum aggregate feeder load $\max_t \sum_{k \in K_f} \sum_{i \in I_k} x_{i,t}$ and whether it saturates $\bar{P}_f^{\text{feeder}}$, (iii) feasibility status of each baseline at each cap value.

F2 result (5 seeds, $n = 500$, $|\mathcal{K}| = 5$, port-count weights $\{6,4,8,5,4\}$). At $\bar{P}_f^{\text{feeder}} = 1$ MW (binding), each decentralised baseline (IsoProp, IsoDem, IsoDual) is feasible in only 3 of 5 seeds — the proportional, demand-weighted, and dual-reallocated splits each leave at least one station too short of headroom on the remaining 2 seeds. CoOpt is feasible in all 5. On the seeds where IsoDem is feasible, the median CoOpt-vs-IsoDem cost edge is $+0.32\%$ [+0.15, +1.10]. At $\bar{P}_f^{\text{feeder}} = 2$ MW (occasionally binding) all four are feasible across all seeds, with a median $+0.88\%$ [+0.61, +1.25] CoOpt-vs-IsoDem edge. At 5 MW (slack) all four collapse to the same solution. This is the qualitative pattern predicted by Proposition 2: coordination matters where capacity is scarce. Median OSQP solve time $\sim$ 2–3 s for the joint problem at $n = 500$.

When coordination pays off, coordination buys little when feeder capacity comfortably exceeds demand. However, it becomes the difference between serving and curtailing the fleet once the shared transformer binds, the regime dense urban clusters (HCMC District 1, the Hanoi Old Quarter) are entering as port counts climb.

Locational scarcity prices. The cap-dual exposure of Proposition 2 is now reported numerically (SmartChargingQP returns per-slot $\nu_{f,t}$ in the solve result; 5-seed median). Figure 3 shows the binding 1 MW feeder cap's shadow price vector across the 24-hour horizon: $\nu_{f,t} = 0$ for $t < 20$ (cap is slack overnight and through midday), then 0.627 kVND/kW at 20:00–21:00 and 1.327 kVND/kW at 22:00–23:00, the feeder binds against the late-evening / early-overnight charging push. Operators reading these prices would prioritise transformer reinforcement on the 22:00–23:00 envelope, where one extra kW of capacity is worth $\approx$ 1.33 kVND of operator cost relaxation.

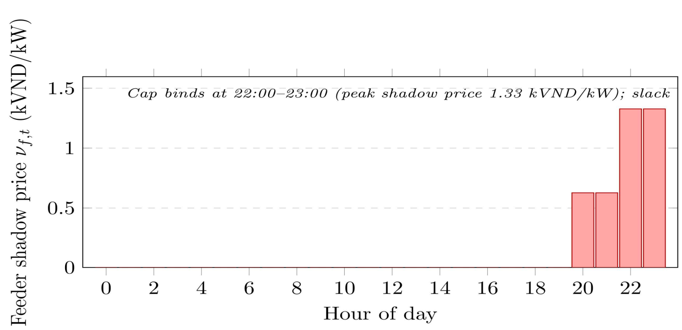

Figure 3. F2 cap-dual vf, t (kVND/kW) per hour at the binding 1MW feeder cap, 5-seed median. The cap is slack from midnight through the early evening and binds only during 20:00-23:00; peak shadow price 1.33 kVND/kW at 22:00-23:00.

3.2.3. F3, tariff sensitivity under the 2025 retail-tariff restructuring

Setting. Fix the F2 instance and sweep the tariff τ_t over four scenarios: 2024-pre-PDP8, 2025-nominal, 2025-flat-midday, and 2025-steep-evening. The scenarios are defined numerically by the bracket vector (off-peak, shoulder, peak) in VND/kWh and the slot-index sets of each bracket (frozen in the artefact archive). The last two scenarios are the structural shifts induced by the 2025 retail-tariff restructuring [2]: flat-midday narrows the shoulder-vs-peak spread between 09:00 and 15:00 to track higher PV share; steep-evening widens the 18:00–21:00 peak relative to off-peak.

We also report the arrival-shift sensitivity sweep, holding the tariff at 2025-nominal, shifting the UrbanEV arrival distribution by $\Delta_a \in \{-2, -1, 0, +1, +2\}$ hours and reporting tariff cost and peak power under CoOpt. We report (i) tariff cost, (ii) peak charging hour, (iii) midday utilisation ratio $\left(\sum_{\text{midday}} \sum_i x_{i,t}\right)/\left(\sum_t \sum_i x_{i,t}\right)$.

F3 result (F2 instance with $\bar{P}_f^{\text{feeder}} = 2$ MW). Tariff sweep, median cost: 2024-pre-PDP8 18.28 M VND; 2025-nominal 17.28 M; 2025-flat-midday 14.91 M (midday utilisation rises from ∼ 20.5% under nominal to ∼ 34.6% as the schedule shifts into the cheaper shoulder hours); 2025-steep-evening 18.54 M. The flat-midday-vs-nominal cost reduction is −14.7%* [−14.8, −13.7] across seeds; steep-evening-vs-nominal is +7.3%. The mechanism behind the flat-midday reduction is shown in Figure 4: *under flat-midday the aggregate station load shifts visibly into the 09:00–15:00 window* (∼ 530 kW sustained) where the bracket-vs-peak spread has been narrowed, while *the late-evening 22:00–23:00 push that defines the 2 MW binding pattern is **reduced*** (1763 vs 2000 kW peak), consistent with the feeder dual reported above.

Arrival-shift sweep (5 seeds, 2 MW modelled feeder cap). Multi-seeding the arrival-shift sweep reveals that the 2 MW cap is genuinely tight even at baseline arrivals: $\Delta_a = 0$ is feasible in only 3 of 5 seeds, $\Delta_a = -1$ h in 3/5, $\Delta_a = -2$ h in 4/5; $\Delta_a \in \{+1, +2\}$ h is infeasible in all 5 seeds. Feasible-seed costs scale monotonically with the shift (medians 17.25, 18.68, 20.13 M VND for $\Delta_a = 0, -1, -2$ h respectively). The infeasibility findings are robust across seeds, but the cliff is shallower than a single-seed snapshot suggests: the 2 MW cap is borderline at $\Delta_a = 0$ rather than comfortably slack, which is itself the operationally interesting message, the modelled feeder is sized at the edge of nominal-day feasibility, and any timing shift exposes it.

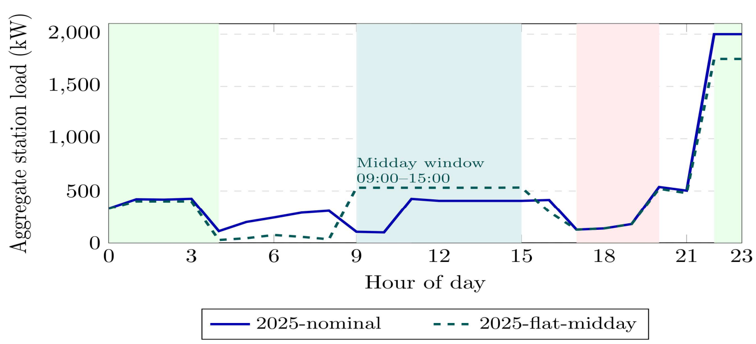

Figure 4. F3 aggregate station load (kW) per hour under 2025-nominal (solid blue) and 2025-flat-midday (dashed teal), 5-seed median. The flat-midday bracket sustains a ∼ 530 kW load through the 09:00–15:00 midday window (teal band), the mechanism behind the −14.7 cost reduction. The late-evening 22:00–23:00 push is also lower under flat-midday (1763 vs 2000 kW), consistent with the feeder-dual pattern in Figure 3.

3.2.4. F4, a solver benchmark

Setting. Fleet sizes $|\mathcal{J}| \in \{50,500,2000\}$; horizon $T \in \{24,48,96\}$. Instances are generated cold (no warm-starting between cells). Solver tolerances are pinned at $\varepsilon_{\text{abs}} = \varepsilon_{\text{rel}} = 10^{-6}$, polish = True, max_iter = 25,000 through CVXPY's solver-options interface. Reference machine: Apple M2 Pro, 16 GB RAM, single thread. Solver wall-clock excludes problem build-time. We report (i) median solver wall-clock time over three seeds per cell, (ii) primal feasibility residual at the returned solution (maximum per-EV energy-delivery error in kWh).

Table 2. F4 result, OSQP with $\varepsilon_{\text{abs}} = \varepsilon_{\text{rel}} = 10^{-6}$, polish=True, max_iter=25,000 (median of 3 seeds; reference machine as stated above).

$\lvert\mathcal{J}\rvert$	T	Median (s)	Max (s)	Median feas resid (kWh)
50	24	0.06	0.06	3.3×10^{-4}
50	48	0.14	0.18	1.9×10^{-3}
50	96	0.33	0.92	3.1×10^{-3}
500	24	2.06	2.62	2.5×10^{-4}
500	48	2.53	2.73	3.5×10^{-3}
500	96	12.32	13.80	8.9×10^{-3}
2000	24	21.50	22.64	7.4×10^{-3}
2000	48	35.57	47.78	1.0×10^{-2}
2000	96	113.3	142.4	6.5×10^{-3}

Discussion. OSQP meets the operational $\leq$ 5s budget at $|\mathcal{J}| \leq 500$, $T \leq 48$; at $|\mathcal{J}| = 2000$, $T = 96$ the observed median wall-clock is $\sim$ 2 minutes, which motivates the Schur-complement decomposition outlined in Section 4. The polish step does not reduce the energy-delivery residual below $\sim 10^{-3}$ kWh per EV at the largest sizes because OSQP's operator-splitting tolerance is exposed to the energy-equality constraint in a way the polish step cannot patch; at $|\mathcal{J}| = 2000$ this is approximately 0.1% of the aggregate daily energy delivered. An interior-point solver or an equality-elimination preprocessing pass would tighten this residual at the cost of additional wall-clock.

4. Conclusion

We formulated Vietnamese multi-station smart charging as a convex quadratic program (1) with a transparent constraint structure: per-EV connection-window bounds, per-EV energy delivery, per-station transformer caps, and feeder-level caps where stations share a primary distribution feeder. The peak-shaving and ramp-smoothing terms are explicit PSD Gram contributions (Proposition 1), and the cap-duals (Proposition 2) yield locational scarcity prices that we used to localise feeder bottlenecks in F2 and exposed numerically via the model. Empirically, across 5 seeds with real UrbanEV-derived arrivals: Smart at $\alpha = 0$ delivers a median 21.6% tariff-cost reduction over ASAP in F1, dropping to $\sim$ 10% at the recommended $\alpha = 10$ in exchange for a roughly two-thirds reduction in station peak power; in F2 at the 1 MW binding feeder cap the decentralised baselines (IsoProp, IsoDem, IsoDual) are feasible in only 3 of 5 seeds while CoOpt is feasible in all 5, with a median cost edge over IsoDem of +0.3–0.9% across binding cap values; the 2025 retail-tariff flat-midday bracket cuts median operator cost by 14.7% in F3; and OSQP solves the $|\mathcal{J}| = 500$ instance in $\sim$ 2 seconds in F4.

Limitations. We highlight five modelling assumptions. First, OpenChargeMap Vietnam coverage is sparse ($\sim$ 6 records), so the F2 five-station feeder is an explicitly modelled scenario informed by the real connector mix rather than a real OCM cluster. Second, no public Vietnamese arrival-trace dataset is available, so arrivals are transplanted from the Shenzhen UrbanEV benchmark; the F3 arrival-shift sweep mitigates the transplant timing effect. Third, transformer caps are calibrated against EVN's per-feeder design heuristics rather than empirical records, and are swept across $\{1,2,5\}$ MW in F2. Fourth, the model is day-ahead deterministic; real-time receding-horizon control and vehicle-to-grid bidirectional flow are out of scope. The F2 coordination claim is a feasibility result paired with a small cost edge (0.3–0.9% median); broader sweeps across cap-splitting heuristics are deferred to follow-up. Fifth, the formulation commits a day-ahead schedule against a fleet manifest $\{a_i, d_i, E_i, \bar{P}_i\}$ that it treats as known exactly;

Solver scaling and the Schur-complement direction. The F4 benchmark shows that the polished OSQP path meets the operational ≤ 5 s budget at $|\mathcal{I}| \leq 500, T \leq 48$ but exceeds it by more than an order of magnitude at $|\mathcal{I}| = 2000, T = 96$. The sparsity recorded in Remark 1 suggests an operator-splitting decomposition along the cap multipliers $\lambda_{k,t}, \nu_{f,t}$ rather than along any single block index of the Hessian.

In the headline $\beta = 0$ regime, the peak-shaving Hessian $\alpha A^\top A$ is block-diagonal in slot index t and the only inter-EV coupling lives inside the per-slot rank-one Gram block $\mathbf{1}_{I_k}\mathbf{1}_{I_k}^\top$ at each station. Dualising the per-station cap (4) and feeder cap (5) with multipliers λ, ν yields a saddle-point reformulation in which, at fixed (λ, ν), the inner problem separates into $|\mathcal{K}|T$ small per-station per-slot QPs of dimension at most $\max_k |I_k|$; this is the regime Lemma 1 solves in closed form when the station is single-EV. A Gauss–Seidel two-pass alternates the inner per-slot QPs (parallel across (k,t)) with a proximal-gradient step on (λ, ν) that enforces the cap residuals. We conjecture that this decomposition can meet the ≤ 5 s budget at $|\mathcal{I}| = 2000, T = 96$, up to the cap-multiplier outer-loop constant and the available parallel-core count on the target machine; the algorithmic specification and complexity analysis are left to follow-up work.

Other future directions. (i) A stochastic QCQP under PV-supply uncertainty with CVaR or variance bounds on tariff cost. (ii) A receding-horizon real-time control loop that re-solves (P) hourly with updated arrival forecasts. (iii) A bidirectional vehicle-to-grid extension that sign-relaxes $x_{i,t}$ and adds an export-tariff term to the objective. All three remain within the convex-QP / convex-QCQP regime; the data-pipeline side is the difficulty rather than the optimisation layer.

References

[1] Prime Minister of Vietnam, Decision 768/QD-TTg approving the revised power development plan VIII, Decision 768/QD-TTg of 15 Apr. 2025; 2025.

[2] Prime Minister of Vietnam, Decision 14/2025/QD-TTg on the retail electricity tariff structure, Decision 14/2025/QD-TTg of 29 May. 2025; 2025.

[3] S. Sojoudi, S. H. Low, "Optimal charging of plug-in hybrid electric vehicles in smart grids," I*n: Proceedings of the IEEE PES general meeting*, Jul. 2011, doi:10.1109/PES.2011.6039236.

[4] L. Gan, U. Topcu, S. H. Low, "Optimal decentralized protocol for electric vehicle charging," *IEEE Transactions on Power Systems*, vol. 28, no. 2, pp. 940–951, May. 2013, doi:10.1109/TPWRS.2012.2210288.

[5] K. Knezović, A. Soroudi, A. Keane, M. Marinelli, "Centralised coordination of EVs charging and PV active power curtailment over multiple aggregators in low voltage networks," *Sustainable Energy, Grids and Networks*, 2021, doi:10.1016/j.segan.2021.100503.

[6] Emily van Huffelen, Roel Brouwer, Marjan van den Akker, "Grid-constrained online scheduling of flexible electric vehicle charging," *Energies*, 2025;18[19]:5063, doi:10.3390/en18195063.

[7] Hongxin Liu, Aiping Pang, Jie Yin, Haixia Yi, Huqun Mu, "Collaborative optimization scheduling strategy for electric vehicle charging stations considering spatiotemporal distribution of different power charging demands," *World Electric Vehicle Journal*, vol. 16, no. 3, p. 176, Mar. 2025;16[3]:176, doi:10.3390/wevj16030176.

[8] Xiang Liao, Ziyu Zheng, Beibei Qian, Haiwei Wang, Dianling Zhan, Junyi Shi, Chaoshun Li, Wei Huang, "Coordinated multi-objective optimization scheduling for electric vehicle swapping station cluster and grid," *iScience*, vol. 28, no. 5, p. 112444, May. 2025, doi:10.1016/j.isci.2025.112444.

[9] Truc Quynh Vu, Hien Minh Ha, Tuan Hai Vu, Electric vehicle charging stations placement optimization in Vietnam using mixed-integer nonlinear programming model [Internet], arXiv:2412.16025; 2024, Available from: https://arxiv.org/abs/2412.16025.

[10] B. K. L. Do, T. H. Nguyen, N. H. Quang, D. Nguyen-Ngoc, L. El Ghaoui, "A digital twin framework for decision-support and optimization of EV charging infrastructure in localized urban systems," *Computers, Environment and Urban Systems*, vol. 127, p. 102422, Jul. 2026, doi: 10.1016/j.compenvurbsys.2026.102422.

[11] P. V. Minh, S. Le Quang, M.-H. Pham, "Technical economic analysis of photovoltaic-powered electric vehicle charging stations under different solar irradiation conditions in Vietnam," *Sustainability*, vol. 13, no. 6, p. 3528, Mar. 2021, doi: 10.3390/su13063528.

[12] T. V. Minh, T. D. T. Kieu, G. V. Pha, T. H. Viet, K. T. Trung, D. V. Ngoc, "Optimal charging scheduling for electric vehicle charging stations with renewable energy integration," Proceedings in Technology Transfer, pp. 174–187, Oct. 2025, doi: 10.1007/978-981-95-1750-3_20.

[13] A. Nguyen, H. Pham, C. Do, "A Cost-Optimization Model for EV Charging Stations Utilizing Solar Energy and Variable Pricing," *Energies*, vol. 18, no. 20, p. 5416, Oct. 2025, doi: 10.3390/en18205416.

[14] OpenChargeMap. Open charge map API. https://openchargemap.io/site/develop/api.

[15] K.-T. Dinh Thi, "Quadratic programming and quadratically constrained quadratic programming: theory, algorithms, and applications," *HPU2 Journal of Science: Natural Sciences and Technology*, vol. 4, no. 03, pp. 80–96, Dec. 2025, doi: 10.56764/hpu2.jos.2025.4.03.80-96.

[16] H. Li, H. Qu, X. Tan, L. You, R. Zhu, and W. Fan, "UrbanEV: An Open Benchmark Dataset for Urban Electric Vehicle Charging Demand Prediction," *Scientific Data*, vol. 12, no. 1, Mar. 2025, doi: 10.1038/s41597-025-04874-4.

[17] J. Frédéric Bonnans, A. Shapiro, "Perturbation Analysis of Optimization Problems," *Springer Nature*, 2000. doi: 10.1007/978-1-4612-1394-9.

[18] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, "OSQP: an operator splitting solver for quadratic programs," *Mathematical Programming Computation*, vol. 12, no. 4, pp. 637–672, Feb. 2020, doi: 10.1007/s12532-020-00179-2.

[19] S. Boyd, "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2010, doi: 10.1561/2200000016.

[20] A. Domahidi, E. Chu, and S. Boyd, "ECOS: An SOCP solver for embedded systems," Jul. 2013, doi: 10.23919/ecc.2013.6669541.

[21] J. Nocedal, S. J. Wright, "Numerical Optimization," *Springer New York*, 2006. doi: 10.1007/978-0-387-40065-5.

[22] B. O'Donoghue, E. Chu, N. Parikh, and S. Boyd, "Conic Optimization via Operator Splitting and Homogeneous Self-Dual Embedding," *Journal of Optimization Theory and Applications*, vol. 169, no. 3, pp. 1042–1068, Feb. 2016, doi: 10.1007/s10957-016-0892-3.

[23] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *The Journal of Risk*, vol. 2, no. 3, pp. 21–41, 2000, doi: 10.21314/JOR.2000.038.

[24] S. Boyd, L. Vandenberghe, Convex optimization, Cambridge University Press; 2004, doi: 10.1017/CBO9780511804441.