Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Verified Bounds on the Charging Gain From Reactive Power at Battery-Buffered Electric Vehicle Fast-Charging Stations

Tien-Khai Nguyen

2430605@student.hcmute.edu.vn

Ho Chi Minh City University of Technology and Engineering

Trong-Nghia Le

Ho Chi Minh City University of Technology and Engineering

Thi-Minh-Chau Le

Hanoi University of Science and Technology

Research Article

Keywords: reactive power, battery energy storage, electric vehicle fast-charging station, Volt-VAR control, convex relaxation, robust feasibility

Posted Date: November 5th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11599477/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Verified Bounds on the Charging Gain From Reactive Power at Battery-Buffered Electric Vehicle Fast-Charging Stations

Tien-Khai Nguyen[1], Trong-Nghia Le[1], and Thi-Minh-Chau Le[2]

[1] Control and Optimization of Smart Energy Systems Laboratory, Department of Electrical and Electronics Engineering, Ho Chi Minh City University of Technology and Engineering, 01 Vo Van Ngan Street, Thu Duc Ward, Ho Chi Minh City, Vietnam
[2] Power Grid and Renewable Energy Laboratory, School of Electrical and Electronic Engineering, Hanoi University of Science and Technology, 01 Dai Co Viet Street, Bach Mai Ward, Ha Noi, Vietnam
Corresponding Author: Tien-Khai Nguyen, Email: 2430605@student.hcmute.edu.vn

Abstract. A battery that supports a fast-charging station has to be recharged through the same feeder, and its converter can supply reactive power while it charges. This paper quantifies how much additional charging that reactive capability admits and how firmly the estimate can be stated. For each control policy, an AC-feasible operating point is paired with an upper bound from an outer linear approximation of the branch-flow relaxation; the dual bound is evaluated in rational arithmetic. A fixed command is then checked over a continuous box of load and source-voltage uncertainty with an invariant voltage enclosure. Four balanced feeders (16, 33, 69, and 141 buses) and controlled changes of R/X ratio, upstream short-circuit level, connectivity, and head-branch rating separate voltage, thermal, and converter limits. At the representative 33-bus location, reactive dispatch raises admissible charging from 250.57 kW to 469.35 kW, and the optimization bracket that remains is 1.1×10^{-4} kW. A Volt-VAR curve tuned for the same state inside the IEEE 1547 Category B setting ranges reaches the same value, so the 124.52-kW shortfall of the default curve is a setting effect. A tuned curve frozen across load levels loses 5.60 kW at the held-out state. In a public 33-bus case, reactive support relieves a binding head-branch limit while the local voltage stays inside the default dead band. All 3340 reported operating points were re-solved with MATPOWER and OpenDSS: voltage magnitudes agree within 7.6×10^{-10} p.u., and no admitted point violates a network limit.

Keywords: reactive power; battery energy storage; electric vehicle fast-charging station; Volt-VAR control; convex relaxation; robust feasibility.

1 Introduction

Extreme fast charging places large power demands at single network connections [1], and the charging load that a feeder can host depends on its voltage, branch-loading, and operating limits [2]. A battery installed beside such a station reduces the peak seen by the feeder, but it must be recharged later through the same connection. The discharge rating of the battery says nothing about how much power the feeder will accept during recharging. The planning question is narrow: if the battery converter supplies reactive power while charging, how much more active power can the feeder admit, and does that gain hold under uncertain load and source voltage?

Reactive-power control is not new. Distributed photovoltaic (PV) studies have examined local and coordinated reactive dispatch [3], [4], local voltage control under limited reactive power [5], and optimized Volt-VAR rules [6]. Charger studies show that reactive operation is technically possible and also report converter and dc-link stresses that battery energy alone does not capture [7], [8]. IEEE 1547-2018 specifies the reactive and voltage-control functions of distributed energy resources (DER), and IEEE 1547.9 applies them to storage [9], [10]. The capability circle of a converter, and the fact that capacitive injection can lift a sagging voltage, are therefore taken as known.

Charging flexibility has been assessed with stochastic AC optimal power flow (OPF) for charging pools [11], with chance-constrained distribution OPF [12], [13], and with robust optimization more generally [14]. Recovery-aware aggregation [15], storage dispatch at charging stations [16], and multi-service storage sizing [17] address the temporal and economic sides of the same problem. Recent work coordinates active and reactive power in vehicle-to-grid stations [18] and combines scheduling with real-time nonlinear OPF on an unbalanced low-voltage feeder [19].

Three weaknesses recur when such studies report a benefit of reactive control. First, the reference policy is often a fixed default curve, so the reported gain mixes the value of reactive hardware with the cost of untuned settings. Second, the optimized value usually comes from a local nonlinear solver with no upper bound, so the reader cannot tell how far it is from the best achievable value. Third, robustness is often judged from a finite set of sampled scenarios, which does not cover the space between samples. A fourth issue arises when results from different test feeders are compared: the feeders differ in voltage level, impedance, load, and source strength at the same time, and the difference is then attributed to topology without a controlled test.

This paper addresses these points for one converter at one location, with hardware held fixed across all policies. Fig. 1 shows how these parts fit together. The contributions are:

1) Each policy result is reported as a bracket. The lower value is an operating point that passes a fresh AC power flow; the upper value is the bound of an outer linear program (LP) of the branch-flow relaxation [20], [21], recomputed from stored matrices in rational arithmetic. The reactive increment is then bracketed without assuming global optimality of any local solver.

2) The gain over unity power factor (UPF) is split into a default-setting loss, a capability effect, and the loss caused by keeping one tuned curve across load states. Snapshot-tuned and frozen Volt-VAR curves inside the IEEE 1547 Category B setting ranges are compared with optimized P–Q dispatch that has the same reactive limit.

3) Controlled interventions on public feeders change the R/X ratio at fixed impedance magnitude, the upstream short-circuit level, the connectivity, and the head-branch rating one at a time. They show that reactive support relieves a binding branch limit while the local voltage is still inside the Volt-VAR dead band.

4) A fixed command is checked over a continuous parameter box with a voltage enclosure based on fixed-point load-flow analysis [22], [23]. The remaining gap is split into enclosure conservatism and search resolution, and failed background states are reported rather than removed.

5) All reported operating points, including 2875 samples inside the uncertainty boxes, are re-solved with MATPOWER [24] and the OpenDSS engine [25], and the conic upper bounds are cross-checked with two independent conic solvers.

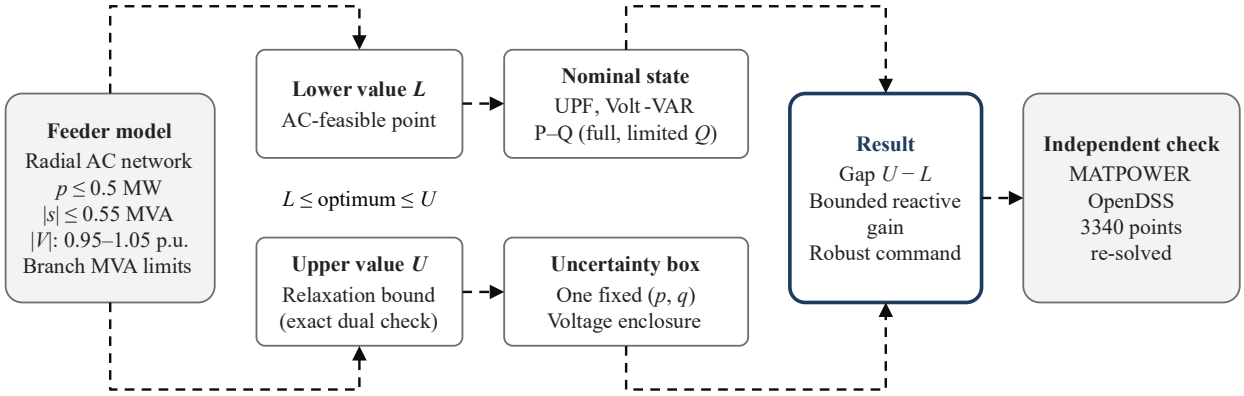


Fig. 1 Assessment workflow. Each policy result is a bracket $[L, U]$; a fixed command is checked over a continuous box; every reported operating point is re-solved with MATPOWER and OpenDSS.

Robust AC feasibility has been studied with convex restrictions [26] and with methods that separate power-flow solvability from engineering limits [27]. Inexact relaxations can produce bounds that do not correspond to an AC operating point [28]; for that reason the objective bound and the AC feasibility evidence are kept separate here.

The 16-bus planning model is taken from an earlier study that placed charging stations on an island grid with mixed-integer optimization and AC power-flow screening [29]. In the present paper, the converter rating and location are fixed and only the control policy changes. Electric-vehicle charging motivates the recharge demand, but the model applies to any bidirectional battery converter; vehicle arrivals, port allocation, and mobility service are not modeled.

2 Network Model and Control Policies

2.1 Network and Converter

Each case is a balanced radial feeder, or a radial forest, with constant-power loads and fixed source-voltage magnitude. One charging converter is placed at bus k. Background demand at bus i is scaled by an active-load factor a and a reactive-ratio factor b. Charging power p is positive when absorbed from the network, and reactive power q is positive for capacitive injection. The net complex demand is

$$s_i = a P_i^d + \mathrm{j}\, ab\, Q_i^d + \mathbf{1}_{\{i=k\}} (p - \mathrm{j}q). \tag{1}$$

The converter rating is 0.5 MW active and 0.55 MVA apparent power for every policy. A point is admissible when the AC power flow has a solution with all bus voltages in 0.95–1.05 p.u. and the apparent power at both ends of each branch within its planning limit $\overline{S}_{ij}$:

$$\begin{aligned} 0 \le p \le 0.5, \qquad & p^2 + q^2 \le 0.55^2, \\ 0.95 \le |V_i| \le 1.05, \qquad & \max\left(\left|S_{ij}^s\right|, \left|S_{ij}^r\right|\right) \le \overline{S}_{ij}. \end{aligned} \tag{2}$$

The MVA limits are planning values. Source import ratings, protection, transformer thermal dynamics, and phase imbalance are not inferred from the available data. In the 16-bus model the two sources feed separate radial components, and the transfer configurations keep the network radial; a source bus left without branches is excluded from voltage monitoring.

Before any charging is admitted, the background state $p = q = 0$ must satisfy the same limits. This rule prevents the converter from being credited with repairing a feeder that is already outside its limits. A state that fails the rule is recorded as background not admitted. A state that fails the rule is recorded as background not admitted and is excluded from the charging results. For an uncertainty box, the background must be admitted over the whole box.

2.2 Policies With Matched Hardware

UPF fixes $q = 0$. Optimized P–Q maximizes p over both signs of q subject to (2). Candidate points are generated by sequential least-squares programming (SLSQP) from four starts: the UPF solution, zero injection, and positive and negative reactive seeds. The selected p is reduced by 10^{-7} MW and the point is checked again with a fresh AC power flow. The multistart only produces candidates; the bound of Section 3 measures how good they are.

The local reference uses the IEEE 1547 Category B default Volt-VAR curve [9], [30] with breakpoints $(0.92, +0.44S)$, $(0.98, 0)$, $(1.02, 0)$, and $(1.08, -0.44S)$, where $S = 0.55$ MVA. The command $q = f(|V_k|)$ is linear between breakpoints and saturates outside them. At each trial p, the local voltage and the reactive command are solved together. Reactive priority is applied: the curve command is kept and p is rejected if the capability circle is violated. An active-power scan followed by bisection gives the largest verified p for this policy. No global certificate is claimed for this nonlinear curve constraint.

The value $0.44S$ is the default reactive magnitude; IEEE 1547 allows settings to be adjusted within stated ranges [9]. Two reactive capabilities are therefore studied: an assumed nameplate value $Q_{\text{cap}} = 0.242$ Mvar, equal to $0.44S$, and the full circle, $Q_{\text{cap}} = 0.55$ Mvar. The optimized P–Q comparator is given the same Q_{cap}; with $Q_{\text{cap}} = 0.242$ Mvar it is labeled P–Q_{lim}. A monotone four-point curve has nine parameters: V_{ref}, four voltage breakpoints, and four reactive values. The two middle reactive values are free, so the dead band may carry a nonzero command. Appendix C lists the ranges and the tuning procedure.

Two tuned curves answer different questions. A snapshot-tuned curve is chosen for one known load state; it is an oracle benchmark and is bounded above by the P–Q relaxation with the same Q_{cap}. A frozen curve maximizes the mean charging over training states $\alpha \in \{0.50, 0.60, 0.70\}$ and is then evaluated, unchanged, at six held-out states. In both cases the reactive law is local, but the active charging level is still screened with the full network model. These are therefore not fully decentralized charging controllers. The comparison is static. Normalized DER settings are applied during battery charging as a modeling choice; conformance and stability testing are outside its scope. Device-specific reactive capability, operator-authorized settings, and dynamic response have to be confirmed for real equipment.

3 Upper Bounds and Continuous-Box Feasibility

3.1 Branch-Flow Relaxation and Brackets

For a directed branch $i \to j$, let P_{ij} and Q_{ij} be the sending-end powers, v_i the squared voltage magnitude, and ℓ_{ij} the squared current magnitude, all in per unit. With branch resistance r_{ij} and reactance x_{ij}, the branch-flow equations are

$$\begin{aligned} P_{ij} &= \sum_{m:j\to m} P_{jm} + p_j + r_{ij}\ell_{ij}, \\ Q_{ij} &= \sum_{m:j\to m} Q_{jm} + q_j + x_{ij}\ell_{ij}, \end{aligned} \tag{3}$$

$$v_j = v_i - 2\left(r_{ij}P_{ij} + x_{ij}Q_{ij}\right) + \left(r_{ij}^2 + x_{ij}^2\right)\ell_{ij}. \tag{4}$$

The equality $P_{ij}^2 + Q_{ij}^2 = v_i\ell_{ij}$ is relaxed to the second-order cone below, which turns the problem into a second-order cone program (SOCP):

$$\left\| \left(2P_{ij},\ 2Q_{ij},\ v_i - \ell_{ij}\right) \right\|_2 \le v_i + \ell_{ij}. \tag{5}$$

The model keeps the converter circle, the apparent-power circles at both branch ends, the squared-voltage limits, and the fixed source voltages. A finite current bound follows from the sending-end limit and the minimum voltage, $\ell_{ij} \le \left(\overline{S}_{ij}/0.95\right)^2$, so every variable has a finite box.

The cones are replaced by an outer polyhedral approximation solved with HiGHS through SciPy [31], [32]. The LP starts from axis halfspaces and adds a supporting halfspace at every cone that is violated. Each cut normal is checked to have norm at most one, and rounding of the cut coefficients is compensated outward over the variable box (Appendix B). An incompletely refined LP is therefore still a relaxation. Exactness of the relaxation is not assumed, and the conditions of [33] are not checked; the LP point is used only for the bound.

Write the encoded LP as $\max c^\top x$ subject to $Ax \le b$, $Ex = d$, and $x^{\mathrm{lo}} \le x \le \overline{x}$. For any $\lambda \ge 0$ and any ν, weak duality with box support gives [34]

$$U = \lambda^\top b + \nu^\top d + \sum_j \max\{\rho_j x_j^{\mathrm{lo}},\ \rho_j \overline{x}_j\}, \quad \rho = c - A^\top \lambda - E^\top \nu. \tag{6}$$

In (6), λ and ν are taken from the LP solver, but exact stationarity is not required: any residual ρ is accounted for by the box term. The bound is computed in exact rational arithmetic from the stored double-precision coefficients and rounded upward. A separate script recomputes it from the stored matrices without calling a solver. The device limit adds $U \le 0.5$ MW. This certifies the arithmetic of the encoded LP; the network data themselves remain model inputs.

Let L and U be the attained AC value and the outer bound for a policy. Besides $U - L$, the reactive increment over UPF satisfies

$$\max\big(0,\ L_{\mathrm{PQ}} - U_{\mathrm{UPF}}\big) \le \Delta p^* \le U_{\mathrm{PQ}} - L_{\mathrm{UPF}}. \tag{7}$$

The bound in (7) uses only the inclusion of the UPF feasible set in the P–Q set. It identifies cases in which the increment is much larger than the remaining optimization uncertainty.

3.2 Fixed Command Over a Continuous Box

A robust command uses one (p, q) for every parameter value in the box and, for the 16-bus model, for all three network configurations. Voltages and flows follow the realized parameters; reactive recourse is not allowed. For uncertainty half-width $\epsilon \in \{0.05, 0.10\}$ and nominal load level α, the box is

$$\begin{gathered} a \in [\alpha(1-\epsilon),\ \alpha(1+\epsilon)], \qquad b \in [1-\epsilon,\ 1+\epsilon], \\ v_0 \in [v_0^c - 0.05\epsilon,\ v_0^c + 0.05\epsilon]. \end{gathered} \tag{8}$$

These are engineering sensitivity ranges, not fitted distributions, so no failure probability is attached to them. When calibrated forecast errors are available, chance-constrained methods [11]–[13] are the better choice; no such data exist for the present cases.

For a cell of (8), the center AC equations are solved first. On a radial network, the voltage satisfies the fixed point $V = T(V; s, v_0) = v_0 \mathbf{1} - Z\,\overline{(s/V)}$, where Z is built from path impedances. Let V^c be the center voltage, $a_v = |V^c|$, d an elementwise bound on $|s - s^c|$, $\bar{s} = |s^c| + d$, $M = |Z|$, and e^c the center residual. For a radius vector $0 < \rho < a_v$, the disks $|V - V^c| \le \rho$ are mapped into themselves for every parameter value in the cell if

$$b_0 + M\left(\frac{\bar{s}\,\rho}{a_v(a_v - \rho)}\right) \le \rho, \qquad b_0 = e^c + \delta v_0 \mathbf{1} + M\left(\frac{d}{a_v}\right), \tag{9}$$

with elementwise products and quotients. The map is a contraction on the disks in the ρ-weighted infinity norm if

$$\kappa = \max_i \frac{[M(\bar{s}\,\rho/(a_v - \rho)^2)]_i}{\rho_i} < 1. \tag{10}$$

Under (9)–(10), a unique voltage solution exists inside the disks for every realization in the cell, and $|V_i| \in \big[a_{v,i} - \rho_i,\ a_{v,i} + \rho_i\big]$. Branch currents are bounded from the same perturbations and multiplied by the endpoint voltage bounds to enclose both branch-end apparent powers. A cell is accepted only if every enclosed quantity meets its limit. Unresolved cells are split, and the accepted cells must cover the whole box. Appendix A gives the derivation.

The command is found in three steps. An optimization over the box vertices gives a candidate (p, q). With q fixed, a scan finds an accepted lower p and a higher tested value, and bisection narrows this interval to at most 10^{-7} MW; every new lower point receives a full box check. Depth budgets of 12 and 18 splitting levels expose the effect of enclosure conservatism. An unresolved upper trial does not prove infeasibility, and because the checker is not shown to be monotone in p, the refined value is reported as an attained point.

The outer relaxation at the box vertices bounds the continuous problem from above, because a command that is feasible everywhere is feasible at each vertex. The total bracket is reported in two parts: the gap between the vertex bound and the vertex candidate, and the derating from the vertex candidate to the enclosed point. The bisection width is reported separately. The enclosure uses residual inflation and positive margins in double-precision arithmetic, instead of directed-rounding interval arithmetic. A failed AC center, an unresolved enclosure, and a fully enclosed box are kept as distinct outcomes.

4 Test Systems and Verification Protocol

Table 1 lists the feeders. The 33- and 69-bus benchmarks originate from [35], [36], and the 141-bus model represents part of the Caracas distribution system [37]. The 33-, 69-, and 141-bus data used in the study are identical to the MATPOWER 8.1 cases *case33bw*, *case69*, and *case141* (largest absolute difference 1.1×10^{-16} in loads and 1.7×10^{-18} p.u. in impedance). The 141-bus loads, given in kVA, are converted at power factor 0.85, and its near-zero-impedance branch 86–87 ($|z| = 6.4 \times 10^{-7}$ p.u.) is kept unchanged. The 16-bus model is the geographically based planning model of [29]. It reproduces the stored voltages of the supplied normal and transfer files to within 4.1×10^{-7} p.u. The main mechanisms are also reproduced on the public feeders.

Table 1 Feeder data and declared operating assumptions.

Model	kV	P/Q (MW/Mvar)	v_0 (p.u.)	Source
16-bus	22.00	11.376 / 2.310	1.00	[29]
33-bus	12.66	3.715 / 2.300	1.02	[35], [24]
69-bus	12.66	3.802 / 2.695	1.02	[36], [24]
141-bus	12.47	11.945 / 7.403	1.02	[37], [24]

Note: Loads are unscaled totals. The 16-bus model has normal and two source-transfer configurations. Public files carry no branch ratings; the 10-MVA limit and $v_0 = 1.02$ p.u. are study settings.

The nominal matrix has six fixed locations per feeder (Table 7 in Appendix D), three load levels $\alpha \in \{0.55, 0.62, 0.72\}$, and three policies, giving 216 cases. Tuned-curve and robust tests use the representative buses 2, 18, 65, and 141 of the 16-, 33-, 69-, and 141-bus models. Controlled experiments use the 33-bus model at bus 18 and $\alpha = 0.62$. Every branch R/X is multiplied by 0.5, 1, or 2 at fixed impedance magnitude. An upstream Thevenin branch with R/X of 0.1 represents an ideal source or a short-circuit level of 250, 100, or 50 MVA behind a fixed internal voltage of 1.02 p.u.; the feeder-head voltage is then free to change. Two virtual reconnections move branch 7–8 to 26–8 and branch 3–23 to 20–23 with the same impedance and loads.

The 16-bus transfer configurations are swept over $\alpha = 0.80$–0.86 at bus 2, where branch loading becomes binding. A uniform branch limit of 6, 10, or 20 MVA is tested on the public feeders. On the public 33-bus model, the charger is placed at bus 2 and the head-branch limit is set to 1.02, 1.05, or 1.10 times its background apparent power at $\alpha = 0.62$, with all other limits at 10 MVA.

Third-party AC verification. All 3340 operating points reported in Sections 5 and the appendices were re-solved outside the study code. Only operating parameters were passed on: network, bus, a, b, v_0, p, q, and the declared intervention. The networks were rebuilt from the MATPOWER case files and from the raw 16-bus data. MATPOWER 8.1 [24] on GNU Octave 8.4 [38] solved each point with Newton–Raphson *runpf* at a mismatch tolerance of 10^{-9} p.u. A tighter tolerance cannot be reached on the 141-bus model because branch 86–87 sets a floating-point mismatch floor near 10^{-10} p.u. The OpenDSS engine (DSS C-API 0.14.5, OpenDSS revision 3723) through OpenDSSDirect.py [25], [39] solved the same points as three-phase circuits in physical units with constant-power loads, $R_0 = R_1$, $X_0 = X_1$, and no shunt capacitance, at a voltage tolerance of 10^{-12} p.u. The points comprise the nominal matrix, all controlled, transfer, rating, tuned-curve, and hourly cases, and a $5 \times 5 \times 5$ grid (including all vertices) inside every enclosed box. The grid is a sampled consistency check; robustness rests on the enclosure. Rejected backgrounds were re-solved to confirm that they violate a limit. For the conic bound, the branch-flow SOCP was re-formulated from the MATPOWER data and solved with Clarabel 0.11.1 through CVXPY 1.9.3 [40], [41], in addition to an earlier MATLAB *coneprog* [42] comparison.

5 Numerical Results and Discussion

5.1 Policy Comparison Across Feeders

At $\alpha = 0.62$, all six locations on the 16- and 141-bus models reach the 500-kW converter limit under every policy (Fig. 2). Reactive control adds nothing at these locations in the nominal state. Voltage limits bind at downstream locations of the 33- and 69-bus models: buses 16 and 18 on the 33-bus feeder and buses 61 and 65 on the 69-bus feeder show the largest differences between policies.

Of the 216 nominal cases, 198 pass the background rule; the 18 rejected cases are the 141-bus states at $\alpha = 0.72$. Each of the 132 admitted UPF and P–Q cases has either the device-cap bound or an outer-LP bound, and the largest bracket $U - L$ is 1.11×10^{-3} kW. At bus 18 of the 33-bus feeder, the P–Q value is 469.348541 kW and its bound is 469.348647 kW. With (7), the reactive increment at this location lies between 218.7803 and 218.7805 kW. The bracket measures optimization uncertainty for the stated model only.

Table 2 Charging at $\alpha = 0.62$ (kW) at the representative buses.

Feeder/bus	UPF	Default	P–Q_{lim}	P–Q	Snapshot	Frozen
16/2	500.00	500.00	500.00	500.00	500.00	500.00
33/18	250.57	344.82	436.25	469.35	469.35	463.75
69/65	333.40	386.03	437.64	461.61	461.61	446.88
141/141	500.00	500.00	500.00	500.00	500.00	500.00

Note: IEEE 1547 Category B default curve. Snapshot and frozen curves use $Q_{\text{cap}} = 0.55$ Mvar; $P - Q_{\text{lim}}$ uses $Q_{\text{cap}} = 0.242$ Mvar.

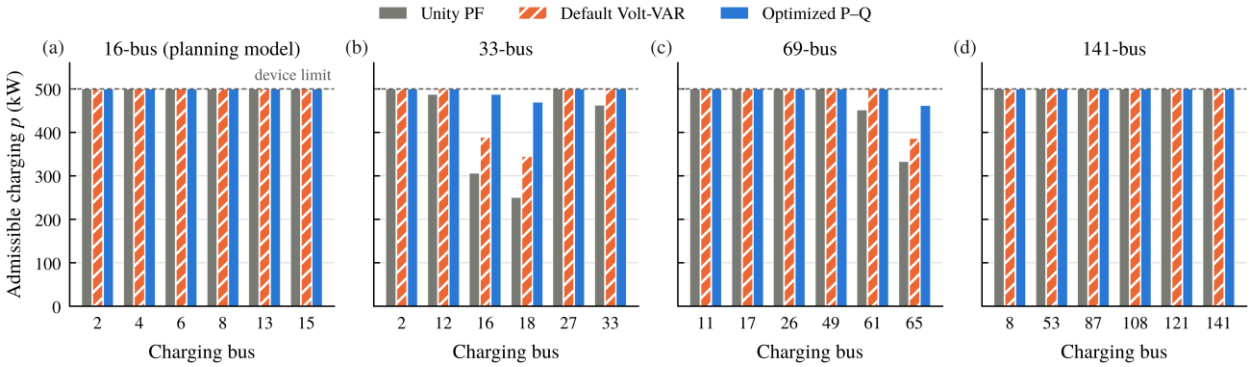


Fig. 2 Admissible charging at $\alpha = 0.62$ for the six fixed locations on each feeder, with 500-kW/550-kVA hardware for every policy.

5.2 Default Setting, Reactive Capability, and Frozen Curves

Table 2 and Fig. 3 separate the parts of the gain. On the 33-bus feeder, P–Q dispatch exceeds UPF by 218.78 kW and the default curve by 124.52 kW. The default-curve shortfall does not survive tuning: a snapshot-tuned curve reaches 469.34864 kW, which is within 9×10^{-6} kW of the P–Q outer bound. On the 69-bus feeder the snapshot-tuned curve is within 4.8×10^{-4} kW of its bound. On the 16- and 141-bus feeders the converter limit closes the comparison. The tuned curve that achieves this is a plateau through the optimized q, which is allowed by the static setting ranges. It serves as an oracle for a known state.

With $Q_{\text{cap}} = 0.242$ Mvar for both the curve and P–Q dispatch, the snapshot-tuned curve again reaches the P–Q value (P–Q_{lim}). On the 33-bus feeder, the 91.43-kW difference between the default curve and P–Q_{lim} is therefore caused by the default settings. The further 33.10 kW comes from the larger assumed reactive capability. On the 69-bus feeder the two parts are 51.61 kW and 23.97 kW. Neither part should be attributed to coordination as such.

The frozen curve, trained on $\alpha \in \{0.50, 0.60, 0.70\}$, gives 463.75 and 446.88 kW at the held-out $\alpha = 0.62$ on the 33- and 69-bus feeders. These values are 5.60 and 14.73 kW below the state-specific P–Q values. Because curve training is nonconvex, the losses include the restriction to one shared curve and any remaining tuning error. The mean of the state-wise P–Q bounds over the training states limits the possible improvement of the frozen tuning to 2.26 kW on the 33-bus feeder and 20.61 kW on the 69-bus feeder (Table 6). On the 141-bus feeder, all training states saturate the converter, so the training objective does not identify a curve and the default setting is kept. Over the five admitted held-out states it averages 424.27 kW against 471.99 kW for state-wise P–Q. A curve trained only on unconstrained states can thus fail when a held-out state becomes voltage-limited.

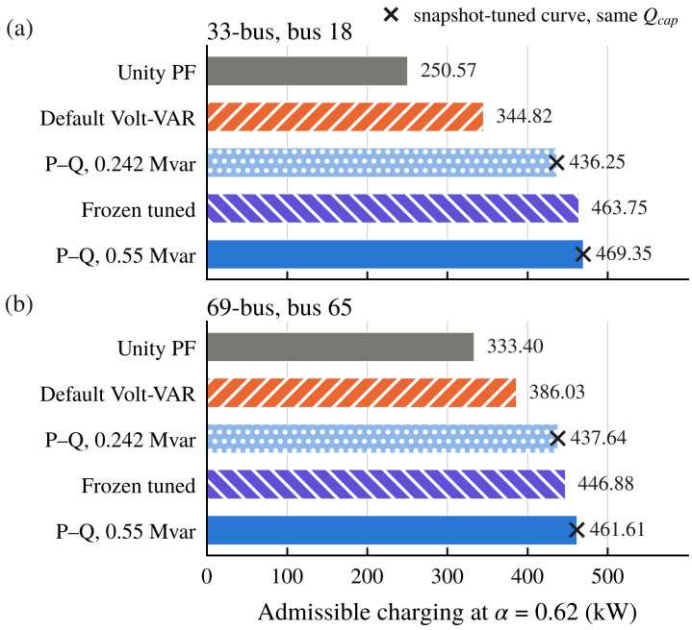


Fig. 3 Decomposition of the reactive benefit at $\alpha = 0.62$. Crosses mark the snapshot-tuned Volt-VAR result with the same Q_{cap} as the bar it lies on.

5.3 Impedance Angle, Source Strength, and Connectivity

Source impedance and feeder R/X interact strongly (Fig. 4). With an ideal source, P–Q charging for R/X multipliers of 0.5, 1, and 2 is 500.00, 469.35, and 375.17 kW, and the increments over UPF are 85.85, 218.78, and 162.02 kW. The increment is not monotonic in R/X. At low R/X, the converter limit caps the P–Q result and hides a large voltage sensitivity; at high R/X, reactive injection moves the voltage less per Mvar.

At the original R/X, reducing the upstream level from an ideal source to 100 MVA lowers UPF charging from 250.57 to 9.02 kW and P–Q charging from 469.35 to 369.38 kW. The reactive increment grows from 218.78 to 360.36 kW while the P–Q capacity itself falls by 99.97 kW. A large relative gain from reactive power can therefore indicate a weak connection rather than a strong one. At 50 MVA the background state violates a limit for all three R/X settings, and no gain is reported.

The two reconnections change the bus-18 P–Q value from 469.35 to 464.10 kW (branch 7–8 moved to 26–8) and to 482.46 kW (branch 3–23 moved to 20–23), with loads and impedances unchanged. These changes are modest. They show sensitivity to connectivity in these two counterfactuals and do not rank topologies. Since the four feeders differ in voltage level, load, impedance, and source settings at once, cross-feeder differences are described in this paper as feeder-dependent, not topology-dependent.

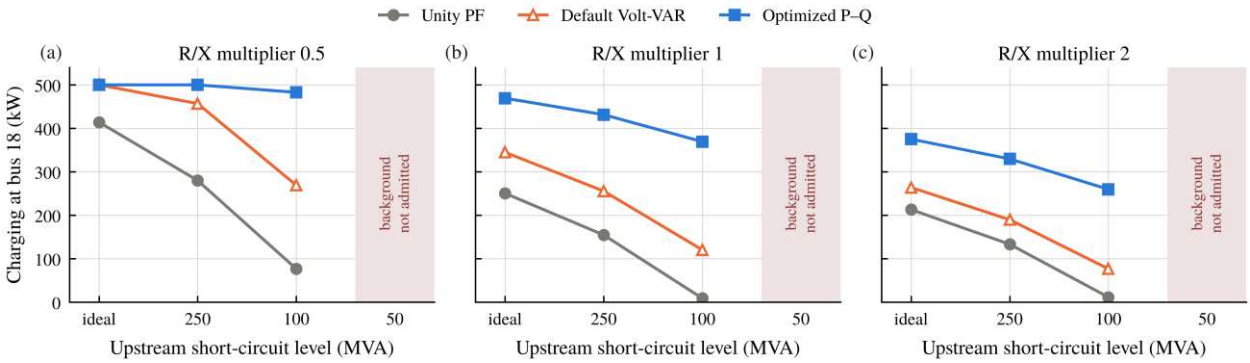


Fig. 4 Controlled R/X and upstream short-circuit interventions on the 33-bus feeder, bus 18, $\alpha = 0.62$. Branch impedance magnitudes are fixed within each R/X setting. At 50 MVA the background state violates a limit for all three R/X settings.

5.4 Branch-Loading Relief Inside the Dead Band

On the public 33-bus feeder, with the head-branch limit set to 2.93335 MVA (1.05 times its background loading), UPF and the default curve both admit 162.32 kW. P–Q dispatch admits 361.41 kW with $q = 0.415$ Mvar (Fig. 5). The bus-2 voltage is 1.0182 p.u., inside the 0.98–1.02 p.u. dead band, and the lowest feeder voltage is 0.9688 p.u., so no voltage limit is active. The head-branch limit binds. Capacitive injection at bus 2 reduces the reactive power drawn through the head branch and releases apparent-power capacity for active charging. The default curve cannot respond because its voltage trigger is never reached.

The effect holds at the other two rating factors: UPF and P–Q admit 65.27 and 291.76 kW at 1.02, and 322.02 and 463.69 kW at 1.10. Snapshot-tuned curves recover the P–Q value at all three factors, because their dead band can carry a nonzero command. The limitation shown here belongs to the default voltage-triggered rule.

The 16-bus transfer configurations show the same mechanism on the planning model. With source SS1 out of service and $\alpha = 0.84$, UPF admits 241.71 kW and P–Q admits 320.16 kW, and the default curve stays at $q = 0$. At $\alpha = 0.86$ the values are 5.06 and 99.91 kW. With SS2 out of service the corresponding values are 247.28 and 324.99 kW at $\alpha = 0.84$. The public intervention reproduces the interpretation without access to the planning files.

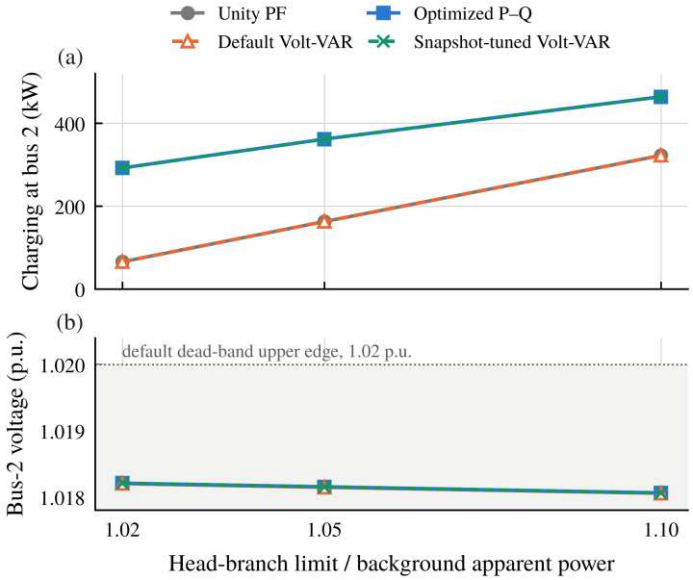


Fig. 5 Public 33-bus head-branch rating intervention, charger at bus 2, $\alpha = 0.62$. The ratings are study assumptions. (a) The default curve coincides with UPF and the snapshot-tuned curve with P–Q. (b) The local voltage stays inside the default dead band for every policy.

5.5 Fixed Commands Over Continuous Boxes

Table 3 and Fig. 6 give the enclosed P–Q commands. The bisection width is at most 10^{-4} kW, so the remaining brackets are not a result of search resolution. For the larger 33-bus box, UPF lies in [85.494, 85.722] kW and P–Q in [386.258, 386.529] kW, and (7) places the robust increment between 300.536 and 301.035 kW. On the 69-bus feeder the robust increment lies between 201.758 and 202.950 kW. Both increments are much larger than the brackets. At $\epsilon = 0.05$ on the 141-bus feeder, the P–Q lower value exceeds the UPF bound by 50.49 kW (Table 8). Converter saturation in the nominal state does not exclude a reactive benefit under uncertainty.

Raising the depth budget from 12 to 18 reduces the 33-bus P–Q bracket at $\epsilon = 0.05$ from 0.481 to 0.120 kW and the 141-bus bracket from 6.702 to 1.677 kW (Fig. 6b). Since the search width is already below 10^{-4} kW, this reduction measures enclosure conservatism. At depth 18 the P–Q covers use 3 cells over the 16-bus configurations, 19 cells for each admitted 33- and 69-bus box, and 145 cells for the admitted 141-bus box. The largest contraction factor is below 0.11. The saved cells were replayed against the inequalities, and their volumes and overlaps were checked in exact rational arithmetic.

At $\epsilon = 0.10$, the 141-bus background violates an AC limit inside the box and is rejected. This is different from an unresolved enclosure, for which failure of the sufficient condition alone proves nothing. The fixed-q refinement also leaves open whether another q or a tighter enclosure would narrow the bracket further.

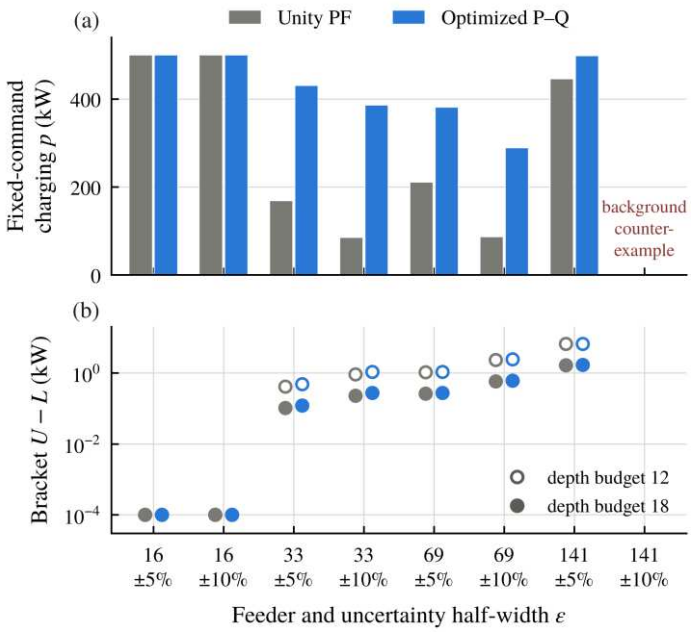


Fig. 6 Fixed commands over the uncertainty box. (a) Enclosed charging for UPF and P–Q at depth budget 18. (b) Bracket $U - L$ at depth budgets 12 and 18 (log scale); the brackets are computational bounds. The 141-bus background fails at $\epsilon = 0.10$.

Table 3 Fixed-command P-Q brackets at $\alpha = 0.62$, depth budget 18.

Feeder/bus	ϵ	Enclosed p (kW)	Bound U (kW)	$U - L$ (kW)
16/2 (3 configs.)	0.05	500.000	500.000	<0.001
16/2 (3 configs.)	0.10	500.000	500.000	<0.001
33/18	0.05	430.648	430.768	0.120
33/18	0.10	386.258	386.529	0.272
69/65	0.05	381.412	381.684	0.273
69/65	0.10	289.008	289.614	0.606
141/141	0.05	498.323	500.000	1.677
141/141	0.10	background counterexample		

5.6 Third-Party AC Verification

Both MATPOWER and OpenDSS converged at all 3340 points (Table 4). The largest voltage difference from the study solver is 7.6×10^{-10} p.u. with MATPOWER, which is consistent with its 10^{-9} mismatch tolerance, and 1.9×10^{-11} p.u. with OpenDSS. None of the 3205 admitted points violates a limit in either tool beyond 3×10^{-12}; such residuals occur only where a limit is binding by design. All 2750 box samples remain feasible with a margin of at least 7.9×10^{-6}, which is consistent with the enclosure being conservative. All ten rejected backgrounds violate a limit in both tools. In the 141-bus box at $\epsilon = 0.10$, both tools find the background violation at the same 10 of the 125 grid points, which reproduces the counterexample.

Table 4 Third-party AC re-solution of all reported operating points.

Point set	N	max $\lvert\Delta\lvert V\rvert\rvert$ (p.u.)		Worst margin
		MATPOWER	OpenDSS	
Nominal matrix	198	5.9×10^{-10}	8.4×10^{-12}	≥ 0
R/X, source, connectivity	54	7.1×10^{-10}	2.3×10^{-12}	≥ 0
16-bus source transfer	24	1.3×10^{-11}	4.3×10^{-12}	≥ 0
Uniform rating, $P - Q_{\text{lim}}$	20	2.5×10^{-10}	6.5×10^{-12}	5.1×10^{-9}
Head-branch rating	12	4.4×10^{-12}	2.2×10^{-12}	3.9×10^{-13}
Tuned Volt-VAR	54	7.6×10^{-10}	1.2×10^{-11}	9.7×10^{-11}
Hourly illustration	93	7.2×10^{-10}	1.2×10^{-11}	≥ 0
Box samples (fixed command)	2750	7.4×10^{-10}	1.9×10^{-11}	7.9×10^{-6}

10

Rejected backgrounds	10	3.1×10^{-12}	8.3×10^{-12}	all < 0
141-bus box, ϵ=0.10, p=q=0	125	3.9×10^{-12}	1.6×10^{-11}	$10 < 0$

Note: $\Delta|V|$: difference from the study solver over all monitored buses. Margin: smallest normalized limit margin from both tools; ≥ 0 means no violation larger than 3×10^{-12} (solver tolerance at binding limits).

Table 5 compares the conic optimum from two solvers with the attained AC value and the certified bound. The Clarabel optimum exceeds L by exactly the 10^{-4}-kW safety reduction applied to every candidate and lies at most 7.5×10^{-4} kW below U. At the four voltage-limited cases, the largest cone slack is below 6×10^{-7} p.u., so the conic solution nearly satisfies the AC equality; the AC points at buses 18 and 65 are therefore within 0.1 W of the relaxation optimum. Where the converter limit binds (16 and 141 buses), the slack is 0.22–0.45 p.u., and the objective agreement there does not indicate an exact relaxation. The *coneprog* values fall up to 4.9×10^{-3} kW below L at its default tolerance, so they serve only as a consistency check.

Table 5 Independent conic solution at $\alpha = 0.62$ (kW).

Case	AC L	Clarabel	*coneprog*	Outer U	Slack
16/UPF	499.999900	500.000000	499.999957	500.000000	0.22
16/PQ	499.999900	500.000000	499.999954	500.000000	0.22
33/UPF	250.568120	250.568220	250.567585	250.568220	3×10^{-9}
33/PQ	469.348541	469.348641	469.346430	469.348647	1×10^{-9}
69/UPF	333.395809	333.395909	333.390954	333.396660	4×10^{-7}
69/PQ	461.614566	461.614666	461.611540	461.615140	5×10^{-7}
141/UPF	499.999900	500.000000	499.999943	500.000000	0.44
141/PQ	499.999900	500.000000	499.999864	500.000000	0.45

Note: Slack: largest $v_i \ell_{ij} - P_{ij}^2 - Q_{ij}^2$ in the Clarabel solution (per unit). *coneprog:* MATLAB R2023b, tolerance 10^{-7}.

5.7 Ratings and Temporal Interpretation

Changing the uniform branch limit from 6 to 20 MVA leaves the 33- and 69-bus results at $\alpha = 0.62$ unchanged, because voltage limits set those values. The 141-bus background fails at 6 MVA and is admitted at 10 and 20 MVA, so its acceptance region cannot be transferred to a physical feeder without branch data. The 16-bus transfer results likewise describe the supplied planning model.

An eight-hour nominal profile (Appendix D) converts the admitted hourly p into battery-side charging energy at a constant efficiency of 0.95. At the 33-bus location, UPF, the default curve, and P–Q give 2.4573, 3.1683, and 3.6881 MWh. These are integrals over an assumed profile and do not constitute a storage sizing result. The 141-bus state at 22:00 fails the background rule, so no complete-window value is reported for that feeder.

6 Conclusion

For a converter that recharges a station battery through a constrained feeder, the charging gain from reactive power was bounded rather than estimated. At the representative 33-bus location, attained AC points and outer bounds evaluated in rational arithmetic place the gain over unity power factor in a 0.2-W interval at 218.78 kW, an 87% increase in admissible charging. Three planning rules follow from the controlled comparisons. First, a default Volt-VAR curve measures the cost of its setting, not the value of reactive hardware: a snapshot-tuned curve inside the IEEE 1547 ranges recovers the full 124.52-kW shortfall, and the default-setting loss (91.43 and 51.61 kW on the 33- and 69-bus feeders) exceeds the gain from raising the reactive limit from 0.242 to 0.55 Mvar (33.10 and 23.97 kW), while keeping one tuned curve across load states costs 5.60 and 14.73 kW at the held-out state. Second, reactive support can raise charging where no voltage limit binds; on a public 33-bus case it relieves a binding head-branch limit while the local voltage stays inside the default dead band, so a voltage-triggered rule gives no response. Third, the relative gain is largest on weak connections, where absolute capacity is smallest, and differences between feeders cannot be assigned to topology without controlled tests. Fixed commands remain feasible over continuous uncertainty boxes with P–Q brackets of 0.12–1.68 kW, and MATPOWER and OpenDSS reproduce all 3340 operating points within 7.6×10^{-10} p.u. These statements hold for balanced constant-power feeders, assumed planning limits, and uncertainty ranges without probabilities; reactive-dependent converter losses, state-of-charge limits, dynamic Volt-VAR stability, and unbalanced three-phase feeders with measured loads remain to be addressed before field use.

Appendix A: Continuous-Box Voltage Enclosure

Consider $T(V; s, v_0) = v_0 \mathbf{1} - Z\,\overline{(s/V)}$ with element-wise division. For a radial forest, every root has the prescribed magnitude and its own phase reference, and disconnected components are handled by the block structure

of Z. Let V^c be a numerical center solution with residual $e^c = |T(V^c; s^c, v_0^c) - V^c|$, $a = |V^c|$, $M = |Z|$, and $|s - s^c| \le d$. The residual is carried explicitly instead of assuming an exact center root.

On the product of closed disks $D = \{V: |V - V^c| \le \rho\}$ with $0 < \rho < a$, every voltage satisfies $|V| \ge a - \rho$, and

$$\left|\frac{1}{V_i} - \frac{1}{V_i^c}\right| \le \frac{\rho_i}{a_i(a_i - \rho_i)}. \tag{11}$$

Writing $s/V - s^c/V^c = (s - s^c)/V^c + s\,(1/V - 1/V^c)$ gives

$$|T(V; s, v_0) - V^c| \le e^c + \delta v_0 \mathbf{1} + M\left(\frac{d}{a} + \frac{(|s^c| + d)\,\rho}{a(a - \rho)}\right). \tag{12}$$

If the right side does not exceed ρ, each parameter-specific map sends D into itself, which is (9). For $V, W \in D$, $|1/V - 1/W| \le |V - W|/(a - \rho)^2$, so the Lipschitz constant of T in the norm $\| z \|_\rho = \max_i |z_i|/\rho_i$ is bounded by κ in (10). If $\kappa < 1$, the Banach fixed-point theorem gives a unique solution in D for every realization in the cell. This is uniqueness inside the enclosure only; other AC solutions elsewhere are not excluded.

For the load parameters with center (a^c, b^c, v_0^c), the componentwise bounds are $\delta P_i = (a^+ - a^-)P_i^d/2$ and $\delta Q_i = \max(|a^- b^- - a^c b^c|, |a^+ b^+ - a^c b^c|)\, Q_i^d$, with $d_i = (\delta P_i^2 + \delta Q_i^2)^{1/2}$ on the feeder power base. All demands and parameter endpoints are nonnegative, so the extremes of the bilinear term ab lie at the stated endpoints. The converter command enters s^c but not d. With H the branch-to-downstream-bus path matrix and $I^c = H\,\overline{(s^c/V^c)}$,

$$\delta I = H\left(\frac{d}{a} + \frac{(|s^c| + d)\,\rho}{a(a - \rho)}\right), \qquad |I| \le |I^c| + \delta I. \tag{13}$$

At each branch end i, the apparent power is bounded by $S_{\text{base}}(a_i + \rho_i)(|I^c| + \delta I)$. Losses enter through the AC center and the nonlinear enclosure, so no lossless DistFlow approximation [35], [36] is involved.

The radius is initialized from the residual and the uncertainty forcing and enlarged until (9) holds or the attempt is rejected. A rejected cell is split along its largest normalized width. A violating center is a counterexample; reaching the depth or cell limit returns *unresolved*. Background admission uses depth 9; command checks use depth budgets 12 and 18 with at most 24 000 cells per trial. Residual forcing is inflated by 10^{-11} p.u., accepted constraints need a margin of at least 10^{-9}, and $\kappa < 1 - 10^{-8}$. The stored leaf boxes, depths, contraction factors, radii, and margins allow the inequalities to be replayed; exact rational volume and overlap checks confirm that the leaves cover the box without overlapping interiors.

Appendix B: Outer-LP Certificate

For the LP of (6), any feasible x satisfies

$$\begin{aligned} c^\top x &= \lambda^\top A x + \nu^\top E x + \rho^\top x \\ &\le \lambda^\top b + \nu^\top d + \sum_j \max\left(\rho_j x_j^{\text{lo}}, \rho_j \overline{x}_j\right). \end{aligned} \tag{14}$$

The argument needs neither an optimal dual vector nor exact stationarity; residuals that would normally be absorbed by bound multipliers enter through the last term. All bounds must be finite and valid for the original feasible set. The code checks $\lambda \ge 0$, evaluates (14) with exact rational copies of the stored double-precision numbers, and rounds upward.

For a cone $\| Bx \|_2 \le \alpha^\top x + t$, any u with $\| u \|_2 \le 1$ gives the valid cut $u^\top B x - \alpha^\top x \le t$. The normal u is scaled conservatively and its squared norm is checked in rational arithmetic. If the exact cut coefficients h are stored as floating values $\hat{h}$, the right side is corrected outward over the variable box:

$$\hat{t} \ge t + \sum_j \max\{(\hat{h}_j - h_j)x_j^{\text{lo}},\ (\hat{h}_j - h_j)\overline{x}_j\}. \tag{15}$$

Cut rounding therefore never removes a feasible point of the encoded cone. Every cone violated by more than 2×10^{-10} receives a cut in each iteration. The loop stops when the cones are satisfied, when the LP objective is within 2×10^{-7} MW of the attained AC value, or after 70 iterations. Early stopping can leave cone residuals; the bound remains valid because of the outer inclusion and (14). Each certificate file contains the sparse matrices, right sides, objective, bounds, multipliers, and LP point. Device-limited cases use the analytical cap.

Appendix C: Volt-VAR Tuning Protocol

Voltages are normalized to nominal, and q is positive for injection. The Category B setting ranges of IEEE 1547-2018 Table 8 [9] are imposed: $V_{\text{ref}} \in [0.95, 1.05]$, $V_2 \in [V_{\text{ref}} - 0.03, V_{\text{ref}}]$, $V_3 \in [V_{\text{ref}}, V_{\text{ref}} + 0.03]$, $V_1 \in [V_{\text{ref}} -$

$0.18, V_2 - 0.02]$, and $V_4 \in [V_3 + 0.02, V_{\text{ref}} + 0.18]$. The reactive values satisfy $0 \le Q_1 \le Q_{\text{cap}}$, $-Q_{\text{cap}} \le Q_4 \le 0$, and $Q_2, Q_3 \in [-Q_{\text{cap}}, Q_{\text{cap}}]$. Monotonicity $Q_1 \ge Q_2 \ge Q_3 \ge Q_4$ and $V_3 - V_2 \ge 10^{-6}$ are added.

For a nonnegative snapshot value q, the settings $(V_{\text{ref}}, V_1, V_2, V_3, V_4) = (1.04, 0.90, 1.04, 1.05, 1.10)$ with $(Q_1, Q_2, Q_3, Q_4) = (q, q, q, 0)$ define a plateau over the observed voltage range. A verified P–Q point on this plateau is feasible for the local policy. Every local-curve point also satisfies the P–Q model, so its value cannot exceed the P–Q bound. An attained plateau point close to that bound is thus a near-best snapshot curve, without solving the nine-parameter design problem globally.

Frozen tuning optimizes the nine parameters jointly with per-state p and q, maximizing their mean over the training states. Six seeds (the default curve and five injection plateaus) start SLSQP with 300 iterations and tolerance 10^{-11}. Each candidate is re-evaluated by a self-consistent AC solution, a 101-point p scan, and 32 bisection steps, with a small positive margin. Training uses $\alpha \in \{0.50, 0.60, 0.70\}$ and evaluation $\alpha \in \{0.49, 0.52, 0.55, 0.62, 0.68, 0.72\}$. On the 141-bus feeder $\alpha = 0.70$ is rejected in training and $\alpha = 0.72$ in evaluation. When several curves reach the converter limit in all training states, the first best curve in the fixed seed order is kept; held-out results are never used to break ties.

Table 6 Frozen Volt-Var curves, $Q_{\text{cap}} = 0.55$ Mvar (kW).

Feeder	Held-out mean				Training objective		
	Adm.	Default	Frozen	P–Q	Frozen	Bound	Gap
16	6/6	500.000	500.000	500.000	500.000	500.000	0.000
33	6/6	370.953	465.691	469.822	465.840	468.104	2.264
69	6/6	395.051	423.789	448.629	428.752	449.364	20.612
141	5/6	424.268	424.268	471.992	500.000	500.000	0.000

Note: Bound: mean of the state-wise P–Q bounds over the training states. Gap: upper limit on the possible improvement of the frozen tuning.

Appendix D: Additional Results

The fixed locations of the nominal matrix are listed in Table 7, and the UPF brackets of the robust study in Table 8. The hourly profile is $\alpha = (0.72, 0.62, 0.55, 0.52, 0.50, 0.49, 0.50, 0.55)$ for hours 22:00 to 05:00. It is an assumed engineering profile, and the hourly commands are optimized separately from the robust commands. The battery-side energy is

$$E_{\text{in}} = 0.95 \sum_{h=1}^{8} p_h \, \Delta t, \qquad \Delta t = 1 \text{ h}, \tag{16}$$

with the results in Table 9. Summing the seven admitted 141-bus hours would not show that the full window is feasible, so that partial sum is not reported. Equation (16) ignores reactive-dependent losses, charge taper, state-of-charge limits, and a terminal energy target.

Table 7 Fixed charger locations in the nominal matrix.

Model	Buses (one-based)
16-bus	2, 4, 6, 8, 13, 15
33-bus	2, 12, 16, 18, 27, 33
69-bus	11, 17, 26, 49, 61, 65
141-bus	8, 53, 87, 108, 121, 141

Table 8 Fixed-command UPF brackets, depth budget 18.

Feeder/bus	ϵ	Enclosed p (kW)	Bound U (kW)	$U - L$ (kW)
16/2 (3 configs.)	0.05	500.000	500.000	<0.001
16/2 (3 configs.)	0.10	500.000	500.000	<0.001
33/18	0.05	168.893	168.995	0.103
33/18	0.10	85.494	85.722	0.228
69/65	0.05	211.037	211.302	0.265
69/65	0.10	86.664	87.250	0.586
141/141	0.05	446.191	447.835	1.645

141/141	0.10	background counterexample

Table 9 Charging energy over the assumed eight-hour profile (MWh).

Feeder/bus	UPF	Default curve	P-Q
16/2	3.8000	3.8000	3.8000
33/18	2.4573	3.1683	3.6881
69/65	3.2341	3.4125	3.6106
141/141	22:00 state not admitted		

Statements and Declarations

Competing interests. The authors have no relevant financial or non-financial interests to disclose.

Author contributions. T.-K.N.: Conceptualization, methodology, software implementation, numerical simulation, formal analysis, and writing—original draft; T.-N.L.: Conceptualization, research direction, methodology review, supervision, and writing—review and editing; T.-M.-C.L.: Supervision, validation, critical review, and writing—review and editing.

Data availability. The 33-, 69- and 141-bus feeder data are available in MATPOWER 8.1. The 16-bus feeder model was developed by the authors in [29]. Its data, together with the scripts, result tables, outer-LP certificates and the MATPOWER/OpenDSS verification files, are available from the corresponding author on reasonable request.

Ethics approval and consent. Not applicable.

Acknowledgments. The authors would like to express their sincere gratitude to Ho Chi Minh City University of Technology and Engineering (HCM-UTE) and Hanoi University of Science and Technology (HUST), Vietnam, for their valuable support and resources provided during this research.

References

1. Tu H, Feng H, Srdic S, Lukic S (2019) Extreme fast charging of electric vehicles: A technology overview. IEEE Transactions on Transportation Electrification 5(4):861–878. https://doi.org/10.1109/TTE.2019.2958709
2. Karmaker AK, Prakash K, Siddique MNI, Hossain MA, Pota H (2024) Electric vehicle hosting capacity analysis: Challenges and solutions. Renewable and Sustainable Energy Reviews 189:113916. https://doi.org/10.1016/j.rser.2023.113916
3. Turitsyn K, Šulc P, Backhaus S, Chertkov M (2011) Options for control of reactive power by distributed photovoltaic generators. Proceedings of the IEEE 99(6):1063–1073. https://doi.org/10.1109/JPROC.2011.2116750
4. Šulc P, Backhaus S, Chertkov M (2014) Optimal distributed control of reactive power via the alternating direction method of multipliers. IEEE Transactions on Energy Conversion 29(4):968–977. https://doi.org/10.1109/TEC.2014.2363196
5. Zhu H, Liu HJ (2016) Fast local voltage control under limited reactive power: Optimality and stability analysis. IEEE Transactions on Power Systems 31(5):3794–3803. https://doi.org/10.1109/TPWRS.2015.2504419
6. Murzakhanov I, Gupta S, Chatzivasileiadis S, Kekatos V (2024) Optimal design of Volt/VAR control rules for inverter-interfaced distributed energy resources. IEEE Transactions on Smart Grid 15(1):312–323. https://doi.org/10.1109/TSG.2023.3279982
7. Kisacikoglu MC, Ozpineci B, Tolbert LM (2013) EV/PHEV bidirectional charger assessment for V2G reactive power operation. IEEE Transactions on Power Electronics 28(12):5717–5727. https://doi.org/10.1109/TPEL.2013.2251007
8. Kesler M, Kisacikoglu MC, Tolbert LM (2014) Vehicle-to-grid reactive power operation using plug-in electric vehicle bidirectional offboard charger. IEEE Transactions on Industrial Electronics 61(12):6778–6784. https://doi.org/10.1109/TIE.2014.2314065
9. IEEE (2018) IEEE standard for interconnection and interoperability of distributed energy resources with associated electric power systems interfaces. IEEE Std 1547-2018. https://doi.org/10.1109/IEEESTD.2018.8332112
10. IEEE (2022) IEEE guide for using IEEE Std 1547 for interconnection of energy storage distributed energy resources with electric power systems. IEEE Std 1547.9-2022. https://doi.org/10.1109/IEEESTD.2022.9849493
11. Giraldo JS, Arias NB, Vergara PP, Vlasiou M, Hoogsteen G, Hurink JL (2023) Estimating risk-aware flexibility areas for electric vehicle charging pools via AC stochastic optimal power flow. Journal of Modern Power Systems and Clean Energy 11:1247–1256. https://doi.org/10.35833/MPCE.2022.000452
12. Dall'Anese E, Baker K, Summers T (2017) Chance-constrained AC optimal power flow for distribution systems with renewables. IEEE Transactions on Power Systems 32(5):3427–3438. https://doi.org/10.1109/TPWRS.2017.2656080

14
13. Baker K, Toomey B (2017) Efficient relaxations for joint chance constrained AC optimal power flow. Electric Power Systems Research 148:230–236. https://doi.org/10.1016/j.epsr.2017.04.001
14. Bertsimas D, Sim M (2004) The price of robustness. Operations Research 52(1):35–53. https://doi.org/10.1287/opre.1030.0065
15. Evans MP, Tindemans SH, Angeli D (2022) Flexibility framework with recovery guarantees for aggregated energy storage devices. IEEE Transactions on Smart Grid 13(5):3519–3531. https://doi.org/10.1109/TSG.2022.3173900
16. Argiolas L, Stecca M, Ramirez-Elizondo LM, Soeiro TB, Bauer P (2022) Optimal battery energy storage dispatch in energy and frequency regulation markets while peak shaving an EV fast charging station. IEEE Open Access Journal of Power and Energy 9:374–385. https://doi.org/10.1109/OAJPE.2022.3198553
17. Arias NB, López JC, Hashemi S, Franco JF, Rider MJ (2021) Multi-objective sizing of battery energy storage systems for stackable grid applications. IEEE Transactions on Smart Grid 12(3):2708–2721. https://doi.org/10.1109/TSG.2020.3042186
18. Xiao F, Li H, Zhang K (2026) Reactive–active power coordination control of grid-forming V2G charging stations for distribution network voltage regulation. World Electric Vehicle Journal 17(5):252. https://doi.org/10.3390/wevj17050252
19. Makhalfih A, Ibrahim IA (2026) A two stage optimal power flow coordination for PVs, batteries, and EVs in unbalanced low voltage distribution networks. IEEE Open Access Journal of Power and Energy 13:641–650. https://doi.org/10.1109/OAJPE.2026.3721623
20. Farivar M, Low SH (2013) Branch flow model: Relaxations and convexification—Part I. IEEE Transactions on Power Systems 28(3):2554–2564. https://doi.org/10.1109/TPWRS.2013.2255317
21. Low SH (2014) Convex relaxation of optimal power flow—Part I: Formulations and equivalence. IEEE Transactions on Control of Network Systems 1(1):15–27. https://doi.org/10.1109/TCNS.2014.2309732
22. Bolognani S, Zampieri S (2016) On the existence and linear approximation of the power flow solution in power distribution networks. IEEE Transactions on Power Systems 31(1):163–172. https://doi.org/10.1109/TPWRS.2015.2395452
23. Bernstein A, Wang C, Dall'Anese E, Le Boudec J-Y, Zhao C (2018) Load-flow in multiphase distribution networks: Existence, uniqueness, non-singularity, and linear models. IEEE Transactions on Power Systems 33(6):5832–5843. https://doi.org/10.1109/TPWRS.2018.2823277
24. Zimmerman RD, Murillo-Sánchez CE, Thomas RJ (2011) MATPOWER: Steady-state operations, planning, and analysis tools for power systems research and education. IEEE Transactions on Power Systems 26(1):12–19. https://doi.org/10.1109/TPWRS.2010.2051168
25. Dugan RC, McDermott TE (2011) An open source platform for collaborating on smart grid research. In: Proceedings of the IEEE PES general meeting, Detroit, MI, USA, pp 1–7. https://doi.org/10.1109/PES.2011.6039829
26. Lee D, Turitsyn K, Molzahn DK, Roald LA (2021) Robust AC optimal power flow with robust convex restriction. IEEE Transactions on Power Systems 36(6):4953–4966. https://doi.org/10.1109/TPWRS.2021.3075925
27. Molzahn DK, Roald LA (2018) Towards an AC optimal power flow algorithm with robust feasibility guarantees. In: Proceedings of the 20th power systems computation conference (PSCC), Dublin, Ireland.
28. Venzke A, Chatzivasileiadis S, Molzahn DK (2020) Inexact convex relaxations of AC optimal power flow problems: Towards AC feasibility. Electric Power Systems Research 187:106480. https://doi.org/10.1016/j.epsr.2020.106480
29. Le TMC, Nguyen TK, Le TN (2026) Grid-impact-aware planning of electric vehicle charging infrastructure for a tourism-intensive island grid using mixed-integer optimization and AC power flow assessment. IEEE Access 14:119449–119466. https://doi.org/10.1109/ACCESS.2026.3718247
30. Huque A (2018) Smart inverters and grid support requirements: IEEE 1547-2018. Presented at the PV Performance Modeling Workshop, 2 May 2018. EPRI. https://pvpmc.sandia.gov/app/uploads/sites/243/2022/10/2-Aminul-Smart-Inverters-and-Grid-Support-Requirements.pdf. Accessed 2 Nov 2026
31. Huangfu Q, Hall JAJ (2018) Parallelizing the dual revised simplex method. Mathematical Programming Computation 10:119–142. https://doi.org/10.1007/s12532-017-0130-5
32. Virtanen P et al (2020) SciPy 1.0: Fundamental algorithms for scientific computing in Python. Nature Methods 17:261–272. https://doi.org/10.1038/s41592-019-0686-2
33. Gan L, Li N, Topcu U, Low SH (2015) Exact convex relaxation of optimal power flow in radial networks. IEEE Transactions on Automatic Control 60(1):72–87. https://doi.org/10.1109/TAC.2014.2332712
34. Boyd S, Vandenberghe L (2004) Convex optimization. Cambridge University Press, Cambridge
35. Baran ME, Wu FF (1989) Network reconfiguration in distribution systems for loss reduction and load balancing. IEEE Transactions on Power Delivery 4(2):1401–1407. https://doi.org/10.1109/61.25627
36. Baran ME, Wu FF (1989) Optimal capacitor placement on radial distribution systems. IEEE Transactions on Power Delivery 4(1):725–734. https://doi.org/10.1109/61.19265
37. Khodr HM, Olsina FG, De Oliveira-De Jesus PM, Yusta JM (2008) Maximum savings approach for location and sizing of capacitors in distribution systems. Electric Power Systems Research 78(7):1192–1203. https://doi.org/10.1016/j.epsr.2007.10.002
38. Eaton JW, Bateman D, Hauberg S, Wehbring R (2023) GNU Octave version 8.4.0 manual: A high-level interactive language for numerical computations. https://www.gnu.org/software/octave/doc/v8.4.0/. Accessed 2 Nov 2026
39. DSS-Extensions (2024) DSS C-API and OpenDSSDirect.py: Alternative implementation of the EPRI OpenDSS engine. https://dss-extensions.org. Accessed 2 Nov 2026
40. Goulart PJ, Chen Y (2024) Clarabel: An interior-point solver for conic programs with quadratic objectives. arXiv preprint arXiv:2405.12762. https://doi.org/10.48550/arXiv.2405.12762
41. Diamond S, Boyd S (2016) CVXPY: A Python-embedded modeling language for convex optimization. Journal of Machine Learning Research 17(83):1–5.
42. MathWorks (n.d.) coneprog: Second-order cone programming solver. Optimization Toolbox documentation. https://www.mathworks.com/help/optim/ug/coneprog.html. Accessed 2 Nov 2026