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.

A diagnostic of voltage infeasibility at low net load and the role of inverter-based grid support

Narciso Moreno-Alfonso

narciso-ma@us.es

Universidad de Sevilla

Alfonso Bachiller-Soler

Universidad de Sevilla

Juan-Carlos Bravo-Rodríguez

Universidad de Sevilla

César Álvarez-Arroyo

Universidad de Sevilla

Research Article

Keywords:

Posted Date: November 3rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11586977/v1

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

Additional Declarations: No competing interests reported.

A diagnostic of voltage infeasibility at low net load and the role of inverter-based grid support

Narciso Moreno-Alfonso[1*], Alfonso Bachiller-Soler[1], Juan Carlos Bravo-Rodríguez[1], César Álvarez-Arroyo[1]

[1*]Departamento de Ingeniería Eléctrica, Escuela Politécnica Superior, Universidad de Sevilla, Calle Virgen de África 7, Sevilla, 41011, Spain.

*Corresponding author(s). E-mail(s): narciso-ma@us.es;

Abstract

At low net load a cost-based unit commitment keeps few synchronous units online, as in the hour of the 2025 Iberian blackout. A violation of the upper voltage limits in a security check is then usually read as an excess of reactive power to be absorbed. This paper examines that reading with two tools. The first is an operability test that poses voltage operability under single generation outages as the minimum-violation problem of a preventive security-constrained optimal power flow. The second is a diagnostic procedure that, with a few more solutions of the same problem, finds which limit carries the violation and whether the state lacks reactive power or has too much. In three test systems and in hours scheduled by a unit commitment, almost every failed state lacked reactive power. The diagnostic exposed how an upper-limit violation can come from a lack: an area without nearby units must be held high to keep a radial load bus within limits, and the end of a lightly loaded cable rises above the ceiling. Capacitors at the radial bus reduced the violation, and reactors at the cable end increased it. With plants at fixed or unity power factor most failures occurred before any outage, and few did with plant voltage control. Batteries that keep charging after their plant trips removed about a quarter of the states operable in voltage control in the largest system. An upper-limit violation is therefore not by itself evidence of excess reactive power, and the diagnostic tells the two situations apart.

Keywords: Inverter-based resources (IBR), Voltage security, Reactive power support, Plant voltage control, Unit commitment, Security-constrained optimal power flow, Co-located battery storage, Low net load, Infeasibility diagnosis

1 Introduction

At low net load, when renewable output covers most of the demand, a cost-based commitment keeps few synchronous units online, and most of the reactive power that holds transmission voltages goes with them. The blackout of the Iberian Peninsula on 28 April 2025, the first large blackout attributed to overvoltage [2, 3], occurred in such an hour [1]. Since then the concern in those hours has been the upper voltage limit [1, 2], and a security check that fails on it suggests absorbing reactive power. This paper shows that such a violation can be the symptom of the opposite need. It proposes an operability test and a diagnostic procedure that tell the two situations apart, and applies them to test systems and to hours scheduled by a unit commitment. Most failures read as overvoltage turn out to be a lack of reactive support, which inverter-based resources (IBR) can supply through plant voltage control and through storage that stops charging when its plant trips.

As the share of IBR grows, such hours become more frequent. The synchronous units they leave offline carry most of the controllable reactive power of a transmission grid: hierarchical voltage control, as practised in France and Italy, holds the voltage of each area through the reactive power of a few control plants [30, 31], and operators have long kept or started units in the areas where voltage support was short [28, 29]. Voltage ranges, however, are set independently of the commitment, and the commitment rules in use protect frequency or system strength [36]. Whether a given commitment at low net load still admits a setting of the voltage controls that keeps every bus within its range before and after any single generation outage therefore has to be checked for each state.

Two opposite expectations exist about how such states fail. Lightly loaded lines produce more reactive power than they consume, so with few synchronous units online the concern has been overvoltage [40]. Yet in the Spanish system the post-contingency violations that required committing units were reported to be low voltages [28], and taking units near the loads out of service increases the electrical distance between generation and load [31]. The Iberian blackout brought the upper limits to the fore: the static single-contingency check was fulfilled shortly before the collapse [1, pp. 422–426], and the investigation recommends both applying the harmonised range of 380 to 420 kV in operation and having converter plants take part in voltage control [1]. Which limit actually binds at low net load, and why, decides which remedy works: absorbing reactive power, or supplying it where it is missing.

Post-contingency voltage limits are standard in a preventive security-constrained optimal power flow (SCOPF), including the discrete taps and shunts of large networks [11, 12, 38], and unit commitment with AC or voltage-stability constraints can choose units that keep enough reactive capability [13, 14, 32]. Planning studies show that a high-renewable generation mix can become infeasible once reactive power is modelled and that the location of reactive sources matters [33–35]. In hours of low load, reactive reserves in absorption become insufficient once generator contingencies are included [15], countermeasures against sustained voltage rise have been compared for a 400 kV system with few units online [16], and the mechanisms of overvoltage have been classified [4, 18]; after the Iberian event, margins to overvoltage-driven collapse and protection-aware screening have been proposed [2, 17, 19]. When a load flow has no

solution, infeasibility can be located and found to come from local reactive deficiency [37]. Closest to this paper, Kuhrmann et al. assess the N-1 static voltage security of a future Nordic system and find that plants in voltage control instead of reactive power control reduce the insecure cases from 17.4 to 4.1% [9]. These works either fix the voltage settings and test a case, optimise a cost for a commitment assumed feasible, or measure loadability margins without single outages. A preventive SCOPF contains the feasibility question but poses it to find a setting of least cost, and infeasibility analysis locates where reactive power is missing in a case without solution. Neither tells an operator whether a violation of the upper limit calls for absorbing reactive power or for supplying it elsewhere, which is what decides the remedy at low net load.

The operability test and the diagnostic procedure answer that question for low-net-load states. Each state is defined by the demand and renewable production of an hour and by the synchronous units that the commitment keeps online. The contributions are:

1. An operability test that poses voltage operability under single generation outages as the minimum-violation problem of a preventive SCOPF, and certifies each positive verdict with load flows at frozen controls.
2. A diagnostic procedure for the states that fail the test. With a few more solutions of the same problem, it finds which limit carries the violation and whether the state lacks reactive power or has too much, and hence the kind of remedy. It needs no knowledge of the structure of the network.
3. A mechanism, exposed by the diagnostic, by which an upper-limit violation comes from a lack of reactive power. An area without nearby units is held high to support a radial load bus, and the end of a lightly loaded cable rises above the limit. Read at face value, the violation points to the opposite remedy.
4. An assessment of inverter-based grid support at low net load. It measures how many failures already occur before any outage, how much of the gain of plant voltage control comes from its reactive capability and how much from its closed loop, and which co-located batteries a requirement on charging after a plant trip must cover.

The results answer four questions:

Q1 Why does a low-net-load state fail: which limit carries the violation, and does the state lack reactive power or have too much?
Q2 Does it fail already before any outage, or only after one?
Q3 Does plant voltage control keep every state that is operable without it, how much does it enlarge the set of operable states, and how much of the gain comes from its reactive capability and how much from its closed loop?
Q4 Does a co-located battery that keeps charging after its plant trips change the answer, and which batteries matter?

Section 2 formulates the test and the diagnostic. Section 3 describes how they are applied: test systems, low-net-load states, models, scope of the assessment, and computation and verification. Section 4 gives the results, Section 5 discusses what they imply for security checks, grid codes and storage, and Section 6 concludes.

2 Voltage operability and its diagnostic

2.1 Voltage operability

In each low-net-load state the operator can still adjust the voltage controls: the set points of machines and plants, the transformer taps and the switchable shunts. The question is whether some adjustment keeps every bus within its limits, before and after any single generation outage. If none does, the state cannot be operated securely with those controls. The commitment, the plant behaviour or the network equipment must then change.

Let $\mathcal{N}$ be the network buses in service, $\mathcal{P}$ the plant-side buses and $\mathcal{E}$ the outages. The vector u holds the voltage controls, within the set $\mathcal{U}$ of available controls, and $V_k^e(u)$ is the voltage of bus k after event e; $e = 0$ is the pre-event state. A state is *voltage-operable* if some $u \in \mathcal{U}$ gives a load-flow solution for every event, with every synchronous machine within its reactive limits, and

$$0.95 \le V_k^0(u) \le 1.05, \qquad 0.95 \le V_k^e(u) \le 1.0875, \qquad k \in \mathcal{N},\ e \in \mathcal{E}, \tag{1}$$

$$0.90 \le V_k^e(u) \le 1.10, \qquad k \in \mathcal{P},\ e \in \{0\} \cup \mathcal{E}. \tag{2}$$

The network limits (Fig. 1) come from the investigation of the 2025 Iberian blackout: the harmonised range of 380 to 420 kV, which it recommends for operation, and the post-outage limit of 435 kV used in the post-contingency check of that system [1, pp. 422, 457]. They are written in per unit so that they apply at any voltage level, and kilovolts in the text are per-unit values times 400 kV. We call 0.95 pu the *floor*, 1.05 pu before the outage the *ceiling* and 1.0875 pu after it the *post-outage limit*. The ceiling and the post-outage limit together are the *upper limits*. A floor of 0.90 pu, the lowest value the European guideline admits [8], is studied as a sensitivity. On the plant side, the range is the band in which inverters are expected to stay connected [7]. The outages are the loss of any one synchronous machine or producing plant, which is the generation part of the single contingencies that planning standards ask to study also off peak [41]. When a unit is offline or trips, its terminal bus and step-up transformer leave the network.

Two sets of available controls bracket what an operator can do. With a *common set point*, every machine and plant regulates to the same voltage, as in a flat voltage schedule. The admissible set points then form an interval. With an *optimal setting*, each machine and plant has its own set point between 0.95 and 1.05 pu, each transformer its own tap between 0.9 and 1.1, and each shunt its own connection. Plants at constant reactive power (Section 3.3) get their own reactive power instead of a set point. This is what a preventive security-constrained optimal power flow would choose [11, 12]. Every state operable with a common set point is operable with an optimal setting, but not the reverse.

2.2 The operability test

To decide whether an optimal setting exists, we look for the setting $u \in \mathcal{U}$ that gives the smallest total violation of the voltage limits over the pre-event state and a subset

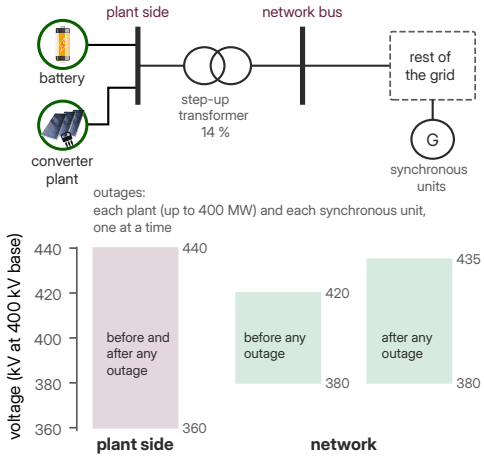


Fig. 1 Voltage limits of the operability test. Top: converter plant, step-up transformer and network bus. Bottom: limits on the network buses before and after an outage, and on the plant side

$\mathcal{E}' \subseteq \mathcal{E}$ of the outages:

$$\phi(\mathcal{E}') = \min_{u \in \mathcal{U}} \sum_{e \in \{0\} \cup \mathcal{E}'} \sum_{k \in \mathcal{N} \cup \mathcal{P}} \Big[\big(V_k^e - \overline{V}_k^e\big)^+ + \big(\underline{V}_k^e - V_k^e\big)^+ \Big], \tag{P}$$

where $(y)^+ = \max(y, 0)$ and $\underline{V}_k^e$, $\overline{V}_k^e$ are the limits of (1)–(2). The voltage magnitudes V_k^e and angles θ_k^e of each event solve its AC load flow with the same controls u:

$$\begin{aligned} P_k^e &= V_k^e \sum_j V_j^e \big(G_{kj}^e \cos\theta_{kj}^e + B_{kj}^e \sin\theta_{kj}^e\big), \\ Q_k^e &= V_k^e \sum_j V_j^e \big(G_{kj}^e \sin\theta_{kj}^e - B_{kj}^e \cos\theta_{kj}^e\big), \end{aligned} \tag{3}$$

Here P_k^e and Q_k^e are the net injections at bus k, $\theta_{kj}^e = \theta_k^e - \theta_j^e$, and $G^e + \mathrm{j}B^e$ is the admittance matrix of the network of event e. The controls u are the set points, the taps, the shunt connections and, for plants at constant reactive power, their reactive output. In each outage the tripped machine or plant is removed and the governors of the units online pick up its active power. Each synchronous machine holds its set point until its reactive output reaches $\underline{Q}_m$ or $\overline{Q}_m$, and then stays at that limit, as in any load flow. Each plant follows its mode, with the droop (4) in voltage control.

The objective adds, in per unit, how far each bus lies outside its limits in each event. A value $\phi \le \tau$, with a tolerance $\tau = 10^{-4}$ pu, means that one setting meets every limit in all those events; Algorithm 1 uses the same tolerance for v_e, the largest violation of an event with the controls frozen. If no setting gives a load-flow solution, (P) has no feasible point; the state then has *no load-flow solution* and is counted apart. The pre-event state is always in (P), so a state that already fails before any outage is rejected without examining the outages. Algorithm 1 builds the subset $\mathcal{E}'$ by adding

Algorithm 1 Operability test of a state

Require: state, outages $\mathcal{E}$, available controls $\mathcal{U}$, tolerance $\tau = 10^{-4}$ pu

$\mathcal{E}' \leftarrow$ all the outages if there are at most 11, otherwise the five with the highest and the five with the lowest voltage at a common set point of 1 pu
repeat
  $u^\star \leftarrow$ minimiser of (P) on $\mathcal{E}'$
  for $e \in \{0\} \cup \mathcal{E}$ **do**
    $v_e \leftarrow$ largest violation of (1)–(2) in the load flow of e with $u^\star$ frozen (∞ without solution)
  end for
  $\mathcal{A} \leftarrow$ the five outages of $\mathcal{E} \setminus \mathcal{E}'$ with the largest $v_e > \tau$; $\mathcal{E}' \leftarrow \mathcal{E}' \cup \mathcal{A}$
until $\mathcal{A} = \emptyset$ or ten rounds
return operable if $\max_e v_e \leq \tau$, and $u^\star$

outages until a load flow with the controls frozen shows no violation, as in contingency filtering for SCOPF [38]. If only events already in $\mathcal{E}'$ exceed τ in that load flow, the limits are tightened by 0.2 kV and the round is repeated once. When the state is not operable, the minimiser of (P) is its *best setting*.

2.3 The diagnostic procedure

When a state is not operable, its best setting shows where the violation is, but not why. The largest violation may appear at a bus far from the problem. A simple case shows how. Take a radial load bus fed from an area without nearby units. Its voltage is roughly the area voltage minus a drop that grows with the reactive power it imports. Without local support, the only way to raise it is to raise the voltage of the whole area. If the area also contains the end of a lightly loaded cable, that bus is already high because of the cable charging. Raising the area then takes the cable end above the ceiling before the radial bus reaches the floor. The largest violation appears at the cable end, but the remedy is reactive support at the radial bus.

The diagnostic (Algorithm 2) uses problem (P) only. It applies to every state that is not operable but has a load-flow solution, and solves (P) on a reduced set of outages $\mathcal{E}_d$. The set is empty if the state fails before any outage. Otherwise it holds the two outages with the largest violations of the upper limits and of the floor, and the five outages with the highest and the five with the lowest voltage at a common set point of 1 pu. If some setting meets every limit on $\mathcal{E}_d$, the outages that violate with that setting frozen are added, as in Algorithm 1, until none does. If a setting then meets every limit only with switched shunts between steps, the shunts are fixed on their steps, as in Algorithm 1, and the diagnostic goes on with them fixed. A state that still admits a setting is left unclassified. The first test finds which limit carries the violation. Problem (P) is solved twice: once with the floor only, the upper limits being raised to 1.30 pu, and once with the upper limits only, the floor being lowered to 0.70 pu. Both values are far from any service range, so this amounts to removing the limit. Table 1 gives the class from which of the two problems has a solution without violation. The second test, the probe, finds whether the state lacks reactive power or has too much. It is

Algorithm 2 Diagnostic procedure for a state that is not operable

Require: state not operable with a load-flow solution, reduced set of outages $\mathcal{E}_d$, tolerance τ

$\phi_0 \leftarrow$ minimum violation (P) on $\mathcal{E}_d$
while $\phi_0 \leq \tau$ and fewer than ten rounds **do**
    add to $\mathcal{E}_d$ the five outages that most violate with that setting frozen, and recompute ϕ_0 ▷ if none outside $\mathcal{E}_d$ violates, tighten the limits by 0.2 kV once
end while
if $\phi_0 \leq \tau$ only with switched shunts between steps, fix them on their steps as in Algorithm 1
if $\phi_0 \leq \tau$ **then return** unclassified
$\phi_F \leftarrow$ minimum violation keeping only the floor, with the upper limits raised to 1.30 pu
$\phi_U \leftarrow$ minimum violation keeping only the upper limits, with the floor lowered to 0.70 pu
class $\leftarrow$ Table 1, from which of ϕ_F and ϕ_U are at most τ
ϕ_+, $\phi_- \leftarrow$ minimum violation with a capacitive and with an inductive susceptance of 1 Mvar at nominal voltage added at every network bus
direction $\leftarrow$ lack if the capacitive susceptance gives the smallest violation, $\phi_+ < \min(\phi_-, \phi_0)$; excess if the inductive one does, $\phi_- < \min(\phi_+, \phi_0)$; otherwise inconclusive
return class, direction and its kind of remedy: supply reactive power for a lack, absorb it for an excess

Table 1 Class of a state that is not operable, from the relaxed problems that admit a setting without violation

Class	Feasible with only the floor ($\phi_F \leq \tau$)	Feasible with only the upper limits ($\phi_U \leq \tau$)
Floor	no	yes
Upper limits	yes	no
Conflict	yes	yes
No equilibrium in the normal range	no	no

a finite-difference sensitivity of the minimum violation to reactive injection: problem (P) with both limits is solved again with 1 Mvar of shunt susceptance added at every network bus, first capacitive and then inductive, and the violation is compared with the original one. Only the probe decides the remedy: a state whose violation lies only on the upper limits can still lack reactive power, as in the case above.

The result of the probe gives the kind of remedy:

- lack of reactive power, if the capacitive susceptance reduces the violation most: supply reactive power in the area, with capacitors, synchronous condensers, plants in voltage control or a unit committed nearby;

- excess of reactive power, if the inductive one does: absorb reactive power, with reactors or plants that absorb;
- inconclusive, if neither reduces it: the need must be located bus by bus.

The diagnostic needs only about five more solutions of (P) per failed state, so any tool that minimises the violation of voltage limits can run it, and it does not use the structure of the network. The probe is a small perturbation, well below the reactive capability of one plant, and only the sign of the change in violation is read. Because it acts on all buses at once, it gives the net direction for the whole state; an area that lacks reactive power and another with too much could cancel out, and locating them bus by bus would need an infeasibility analysis [37]. A finite difference is used rather than the multipliers of the voltage constraints because the machines switching at their reactive limits and the shunt steps make the minimum violation non-smooth, so its local derivatives are not reliable.

3 Assessment procedure

3.1 Test systems

To answer Q1 to Q4 we need systems where units can be taken offline one at a time in a known order, where plants and storage sit where a high-renewable system would place them, and where the hours an operator would schedule can be examined. We use three open test systems, each for a different reason (Table 2). The 9-bus system of Anderson and Fouad, a reduced model of the Western Systems Coordinating Council (WSCC) system generally known as the IEEE 9-bus system [20], is a minimal case to see the effect of plant mode and battery. The IEEE 39-bus system [21] is a 345 kV meshed network with long lines and a reference case for voltage stability; it tests whether the findings hold in a network of realistic structure. The Reliability Test System of the Grid Modernization Laboratory Consortium, RTS-GMLC [22], has 73 buses in three areas, radial load buses and a cable, unit costs, and a year of hourly demand and renewable resource. It is one of the few open systems that allow both a decommitment sequence and a unit commitment on the same network. As in the published case, a synchronous unit of the RTS-GMLC groups the machines of one type at a bus.

3.2 Low-net-load states

Two kinds of state answer different questions. *Decommitment sequences* isolate the effect of the commitment. A sequence starts with all units online at their technical minimum and takes them offline one at a time, the most expensive first, down to the smallest fleet that can serve the hour. Consecutive states therefore differ by one unit. The sequences are nodal: each step removes a given unit at its bus. Four scenarios combine a moderate and a high renewable capacity with the published network and with a weakened one (Table 2). The weakened network lengthens the electrical distance between generation and load. It multiplies by 1.6 the resistance and reactance, and divides by 1.6 the charging susceptance and the thermal rating, of the branches on the shortest electrical paths between the machines and the two load buses with the largest demand and no machine. Step-up transformers are not changed. This affects 4 of the

9 branches of the 9-bus system, 29 of 46 in the 39-bus system and 72 of 120 in the RTS-GMLC. The sequences use the same synthetic design day in every system and scenario, with a clear-sky solar profile, a wind profile and a parametric demand curve. From that day we take the hour of least demand and the hour of largest renewable resource relative to demand; they are the same in every scenario. Both are hours of *low net load*, the demand minus the renewable production: the demand is 0.60 and 0.87 of the peak, and at the hour of largest resource the renewable production reaches 1.6 times the demand. Storage charges up to its rating and the rest of the surplus is curtailed. A DC linear programme then curtails plants, reduces charging or raises units until every branch is within its rating. Consecutive states of a sequence are not independent, so results are also given sequence by sequence (Fig. 4).

Table 2 Test systems and low-net-load states

	IEEE 9-bus	IEEE 39-bus	RTS-GMLC
What it shows	Plant mode and battery in a minimal case	Meshed transmission network	Multi-area system with radial buses, and what a unit commitment produces
Network	9 buses, 230 kV	39 buses, 345 kV	73 buses, 138 and 230 kV
Synchronous units	3	10	49, with 96 machines, 3 of them synchronous condensers
Converter plants	5, one per bus	20 at 14 load buses	30 at 21 load buses; in the scheduled hours, 21 at their published buses
Co-located batteries	2	10	15; also 15 in the scheduled hours
Converter rating, plants and batteries, E1/E3 and E2/E4	425 and 866 MW	8443 and 17 199 MW	11 543 and 23 513 MW in the sequences
Battery rating	142 MW	2814 MW	3848 MW in the sequences; 3686 MW in the scheduled hours
Switchable shunts	none	none	reactors of 100 Mvar at buses 106, 206 and 306
Units online in the sequences	1–3	1–10	1–49
Units online in the scheduled hours	–	–	9–22
Scenarios	E1 and E2: renewable capacity of 0.9 and 2.3 times the peak demand on the published network; E3 and E4: the same on the weakened network of Section 3.2		
Hours of each scenario	Largest renewable resource and least demand, one decommitment sequence each		

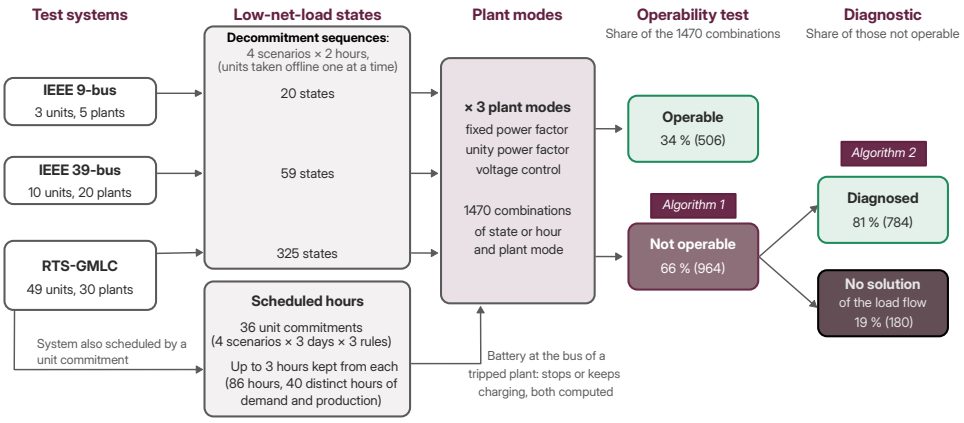


Fig. 2 Design of the assessment: test systems, low-net-load states, plant modes, operability test (Algorithm 1) and diagnostic procedure (Algorithm 2). Shares of all the combinations and, for the diagnostic, of those not operable

Scheduled hours show what an operator would face. Only the RTS-GMLC has the unit costs and hourly series that a unit commitment needs. The day-ahead commitment keeps the published plant siting, limits branch flows under single contingencies, and limits the rate of change of frequency and the frequency nadir after the largest outage. It is solved for three days of the published year in each scenario: the day with the largest ratio of renewable resource to demand, the day with the largest hourly rise of net load, and the day with the median ratio. It uses three commitment rules: no rule beyond the frequency limits; at least eight units of 155 MW or more online, a rule studied for the Irish system [23]; and a frequency constraint evaluated at the operating point. The rules also show whether a minimum number of large units online secures voltage operability. The commitment is zonal: it decides how many units of each type are online in each area, and these are shared among the buses of the area in proportion to the units each bus has. From each commitment we keep up to three hours: the hour with fewest units online, the hour of least demand and the hour with the largest share of renewable production; repeated hours count once. The demand in those hours is 0.36 to 0.58 of the peak. Different rules often pick the same hour with a different commitment, so some scheduled hours share demand and production. Fig. 2 gives the counts. In what follows, a state is a step of a sequence or a scheduled hour, and a *combination* is a state with a plant mode. The four *sets of states* are the sequences of each system and the scheduled hours of the RTS-GMLC.

3.3 Plant, machine and outage models

The plants are sited by rule in the 39-bus system, which has no published renewable siting, and in the RTS-GMLC sequences, whose published siting cannot export its production at low net load. Solar plants go to the load buses without a machine whose demand is above the median of those buses, with a rating in proportion to that demand.

The same number of wind plants of equal rating go to the eligible buses electrically farthest from any machine, and a battery sits at each solar plant. In the 9-bus system the five plants sit at buses 4 to 8. The scheduled hours keep the published siting of the RTS-GMLC. The renewable capacity is 0.9 or 2.3 times the peak demand and the battery rating 0.45 times the peak demand (Table 2).

Each converter plant connects through step-up transformers with a total reactance of 14% on the converter rating, 6% for the collector transformers and 8% for the main transformer, within the range of the guide of the Western Electricity Coordinating Council [7]. The plant injects at the plant-side bus, where the inverter protections measure. Plants run in one of three modes. At *fixed power factor* they absorb at 0.98 inductive, a common grid-code setting and the one used by the Iberian plants in April 2025 [1, pp. 91, 96]; this pulls the voltages down before any outage. At *unity power factor* they exchange no reactive power. In *voltage control* each plant regulates the voltage of its network bus with a droop,

$$Q = Q_0 - \frac{Q_{\text{cap}}}{\sigma}\,(V - V_{\text{ref}}), \qquad |Q| \le Q_{\text{cap}} = 0.30\,S_{\text{conv}}, \tag{4}$$

where Q_0 is the reactive power at the reference voltage V_{ref} and σ the slope. The slope is 3.33%, within the adjustable range of 2 to 7% of the European network code [6]; 2 and 7% are studied as sensitivities. The capability Q_{cap} is 30% of the converter rating at any production, as in the Spanish voltage-control procedure revised in 2026 [5]. With a slope of 3.33%, a plant gives its full capability when its bus voltage moves 3.33%, about 13 kV, from its reference.

Two comparators separate what voltage control gains from its capability and from its closed loop. Plants at *constant reactive power* inject $Q = \alpha\,Q_{\text{cap}}$, with α between -1 and 1 chosen by the optimal setting and kept after the outage; unity power factor is $\alpha = 0$. Voltage control is also computed with capabilities of 15 and 45% of the converter rating. The 15% case represents plants that offer less capability when they produce little.

Synchronous machines regulate their terminal voltage and switch to fixed reactive output at a limit. In the 39-bus system the published minima are mostly zero or positive, so every machine is allowed to absorb down to 0.30 times its active power rating. Each synchronous machine trips on its own, even where a unit groups several machines, and so does each producing plant. The active power lost in a plant outage is limited to 400 MW in the 39-bus system and the RTS-GMLC, the largest machine of the RTS-GMLC, and to 247.5 MW in the 9-bus system, its largest machine, even if the plant produces more. A tripped plant takes its share of the converter capability of its bus, and the governors of the units online pick up the power lost in proportion to their rating. A battery at the bus of the tripped plant either *stops charging*, as when the plant controller or a protection acts on it, or *keeps charging* on its market schedule. Which one happens depends on its control and protection, so both cases are computed for every state.

3.4 Scope of the assessment

The assessment rests on four modelling choices:

- Each outage is assessed at the quasi-static equilibrium reached seconds after it, before tap changers or relay timers act, as in static security assessment [4, 25]. Phenomena within the protection window [17] and the later action of tap changers [24] are not covered. Every plant mode is assessed with the same model.
- Only generation outages are studied: each synchronous machine and each producing plant, with the limit on the power lost of Section 3.3. Every comparison uses this same set.
- The floor is 0.95 pu before and after the outage, while the upper limit rises to the post-outage limit after it. The setting is preventive: no corrective action follows the outage. A floor of 0.90 pu before and after the outage is studied as a sensitivity; the plant side keeps its range of 0.90 to 1.10 pu.
- The classes and the probe of the diagnostic are computed with the switched shunts treated as continuous. A test with only one limit could therefore admit a setting with shunts between steps that it would not admit with the shunts on their steps.

3.5 Computation and verification

The common set point is swept from 380 to 420 kV in steps of 5 kV, refined to 1.25 kV where the interval is narrower than a step, with a load flow of every event at each set point. The sweeps drawn in Figs. 9 and 11 use a step of 2.5 kV. Problem (P) is solved with IPOPT [26] through CasADi [27] from four starting points, and the best solution is kept. The violations $(y)^+$ are written as non-negative slack variables. The switching of each machine at its reactive limits is written as complementarity conditions, relaxed by a small tolerance. In both algorithms the switched shunts are first treated as continuous between 0 and 1, which can only widen the feasible region, and are then fixed on their steps: rounded and, if the rounding fails, with each of the three closest to 0.5 switched in turn. In 508 of the 784 combinations diagnosed, Algorithm 2 was run from two starting points instead of four, and the probe was applied only to the classes in which at least one limit can be met on its own.

IPOPT is a local solver, and overvoltage limits can split the feasible region into separate parts [10]. For this reason, every state found operable is confirmed by load flows with the controls frozen, so it is operable whatever the solver. A state found not operable rests on the best of four starting points, so the counts of operable states are lower bounds. With the floor of 0.90 pu, the diagnostic found one such state, with a setting that the load flows with the controls frozen confirm; it is left in the counts as Algorithm 1 gave it. A lower floor can only add admissible settings, so every state operable with a floor of 0.95 pu must stay operable with 0.90 pu; all of them do.

The diagnostic was validated on states with a known cause. We took the 118 operable states of the 9-bus and 39-bus systems at unity power factor and in voltage control. To create an excess, the charging susceptance of every line was multiplied by 2 to 8 until the pre-event state had no admissible setting; to create a lack, the reactive demand of every load was multiplied in the same way. Some states lost their load-flow solution before reaching that point, and the rest were diagnosed. With the created

excess, the probe found excess in all 93 states diagnosed. With the created lack, it found lack in 81 of 96. The other 15 were conflicts of a few kilovolts in the 39-bus system in voltage control, where the net direction over all buses came out the other way.

The remedy indicated by the diagnostic was assessed on the 20 *remedy hours*: the scheduled hours at unity power factor with no admissible setting before any outage and with the largest violation on the upper limits. In each, problem (P) for the pre-event state was solved again twice: with 50 Mvar of capacitors at bus 307, the radial load bus of area 3 of the RTS-GMLC, and with 50 Mvar of reactors at bus 306, the end of the 306–310 cable. The reactors are the remedy that the largest violation suggests.

4 Results

Section 4.1 gives what the diagnostic finds in the failed states (Q1). Section 4.2 explains the mechanism behind upper-limit violations. Section 4.3 shows how many failures occur before any outage (Q2). Sections 4.4 and 4.5 show how plant voltage control and co-located storage change the results (Q3 and Q4). Section 4.6 summarises the answers.

4.1 Diagnosis of the failed states

Almost every failed state lacks reactive power, whatever limit carries its violation (Fig. 3a). Fig. 2 shows how many combinations reach each step of the diagnostic. The probe was run in every class; in the class without equilibrium in the normal range it covered 54 of its 99 combinations (Section 3.5), and 52 of them lacked reactive power. In that class the relaxations already show that no setting meets the floor or the upper limits, even on their own.

In 99% of the combinations probed, a capacitor at every network bus reduces the violation more than a reactor. Only 2 of the 739 combinations probed show an excess, and 2 are inconclusive. The lack appears in every class. Conflicts are the largest class, with 63% of the combinations diagnosed: the floor and the upper limits can each be met, but not both at once. Fewer fail on the floor alone, and only a few on the upper limits alone. With a floor of 0.90 pu the result is the same: 99% of the combinations probed lack reactive power.

This contradicts the usual reading of the largest violation. With plants at fixed or unity power factor, the largest violation of the best setting often lies on the upper limits, which suggests absorbing reactive power. Yet in these combinations a capacitor reduced the violation in 98% of those probed, and a reactor increased it in 93%. With plants in voltage control the largest violation almost never lies on the upper limits. Fig. 4 shows the RTS-GMLC sequences state by state; with 49 units they have enough steps to show where the failures start, which the 9-bus and 39-bus systems, with 3 and 10 units, do not. Almost every failure is a lack of reactive support, and the upper limits alone bind only in a few states with few units online.

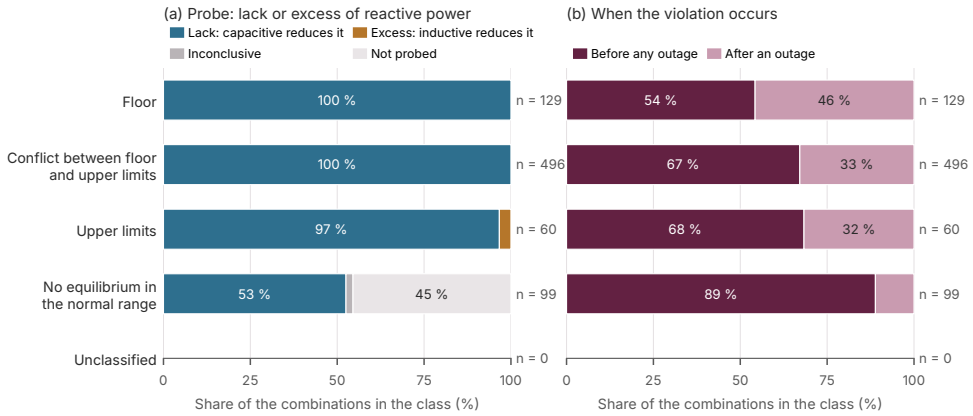


Fig. 3 Diagnostic of the combinations that are not operable, by class, over the four sets of states and the three plant modes: (a) result of the probe; (b) violation before or after an outage. Bars give shares of each class; n, combinations in the class

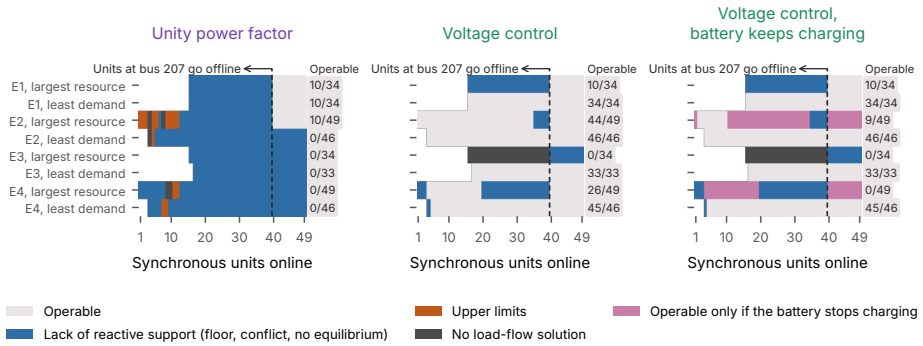


Fig. 4 Limit that carries the violation in each state of the RTS-GMLC decommitment sequences, at unity power factor and in voltage control with the battery that stops or keeps charging. Rows: sequences by scenario and hour. Right: operable states over states. Dashed line: units at bus 207 go offline

4.2 Upper-limit violations from a lack of reactive power

In the scheduled hours the largest violation lies mostly on the upper limits. The diagnostic traces it to a lack of support in the south of area 3 of the RTS-GMLC (Fig. 5a). Bus 307 is a radial load bus fed only through bus 308. To keep its voltage up, the best setting raises the whole area, and bus 306, at the end of the 306–310 cable, goes above the ceiling. In the 5 remedy hours whose best setting gives a load flow with the controls frozen, the highest voltage is always at bus 306.

Fig. 5b shows one remedy hour: hour 12 of the high net-load ramp day of E1, scheduled with the frequency constraint at the operating point, with 14 units online and plants at unity power factor. Without local support, bus 307 is below the floor

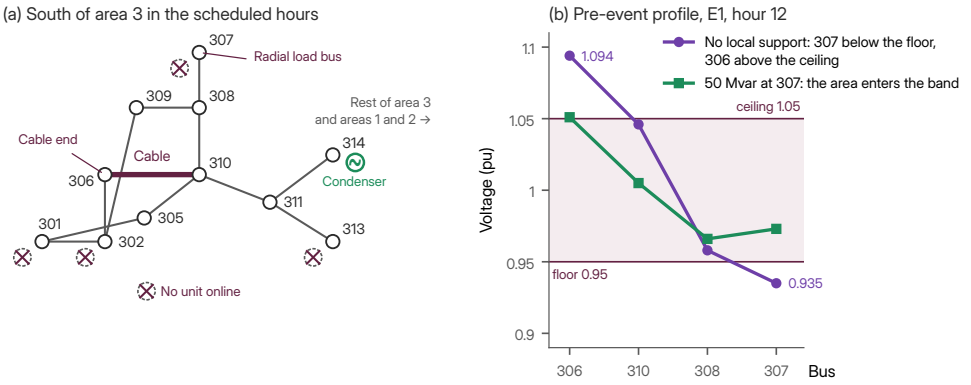


Fig. 5 Lack of reactive power in the south of area 3 of the RTS-GMLC: (a) network around the radial load bus 307 and the 306–310 cable; (b) pre-event voltage profile along buses 306, 310, 308 and 307 in one remedy hour, without and with 50 Mvar of capacitors at bus 307

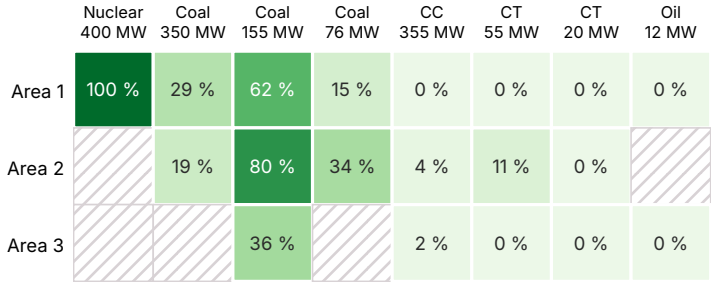


Fig. 6 Share of the 864 committed hours with at least one unit of each type online, by area of the RTS-GMLC. CC: combined cycle; CT: combustion turbine. Hatched: no unit of that type in the area

and bus 306 above the ceiling; with 50 Mvar of capacitors at bus 307 the whole area enters the band. Capacitors at bus 307 reduced the violation in every remedy hour that kept a solution and removed it before the outage in 4 of the 20 hours. Reactors at bus 306, the remedy that the largest violation suggests, increased the violation or left the hour without a solution in all of them. The high voltage at the cable end was a symptom of the lack at the radial load bus.

The lack comes from the commitment (Fig. 6). In the south of area 3, buses 301, 302 and 307 have only combustion turbines and bus 313 a combined cycle. With abundant renewable energy, a cost-based commitment leaves these units offline in almost every hour, while areas 1 and 2 keep coal and nuclear units online. Since these types are offline in the whole area, no unit is ever online at bus 307, whatever the allotment among buses. A rule on the number of units does not help: with at least eight large units online, 65% of the scheduled hours are operable in voltage control, against 77% without any rule.

In the sequences, location shows up as a single unit deciding operability (Fig. 4). In E1, with moderate renewable capacity on the published network, the sequence of largest resource loses operability in voltage control when the combustion turbines at bus 207 go offline. Bus 207 is the radial load bus of area 2, fed only through bus 208. Below that step, the outage of the plant at bus 208 or 308 takes the lowest voltage below the floor, and voltage control gains nothing over unity power factor. The sequence of least demand of the same scenario stays operable down to its smallest fleet. On the weakened network, E3, the sequence of largest resource is not operable even with all units online. With high renewable capacity, E2 and E4, the plants have larger converters and therefore more reactive capability, and the sequence keeps operable states with very few units online.

4.3 Failures before any outage

Without plant voltage control, most failures occur before any outage. At fixed or unity power factor, 60% of the combinations that are not operable already fail in the pre-event state; in voltage control only 15% do, all of them in the scheduled hours. Fig. 3b gives the same split by class. The RTS-GMLC sequences show the contrast most clearly. At unity power factor 71% of the failed states fail before the outage, and at fixed power factor 60%, because another 34% have no load-flow solution. In voltage control none fails before the outage, and the post-outage check becomes the limit. This answers Q2.

4.4 Plant voltage control

Plant voltage control is the remedy the diagnostic points to among inverter-based resources. It made no operable state inoperable: in every set of states and with both battery behaviours, every state operable at fixed or unity power factor is also operable in voltage control. It raised the share of operable states from 9% at unity power factor to 73% in the RTS-GMLC sequences, and from 23% to 72% in the scheduled hours (Figs. 7 and 8, Table 3). In the sequences a common set point already gives most of the gain, 95% of the states operable with the optimal setting. The scheduled hours need individual set points, taps and shunts; there a common set point reaches only 35% of them.

Most of the gain comes from the reactive capability; the closed loop adds the rest (Table 3). Plants holding an optimised constant reactive power with the same capability make fewer states operable than plants in voltage control. The difference is the states where no fixed reactive output suits both the pre-event state and every outage. Voltage control with half the capability, 15%, does at least as well as constant reactive power with 30%: better in the 39-bus system, the RTS-GMLC sequences and the scheduled hours, and equal in the 9-bus system, where every state is operable. Raising the capability to 45% adds little, and in the 9-bus and 39-bus systems the capability does not change the result.

Fig. 9 shows how voltage control acts, by sweeping the common set point in four states of the RTS-GMLC. At unity power factor, the lack shows at the bottom of the voltage profile. One state is operable only near 404 kV, where the floor after the outage

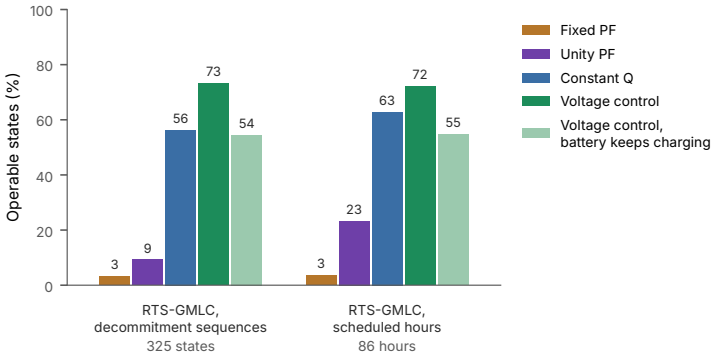


Fig. 7 Voltage-operable states of the RTS-GMLC with the optimal setting, by plant mode, in its decommitment sequences and scheduled hours. Constant Q: optimised constant reactive power within a capability of 30%. Unless stated, the battery stops charging when its plant trips

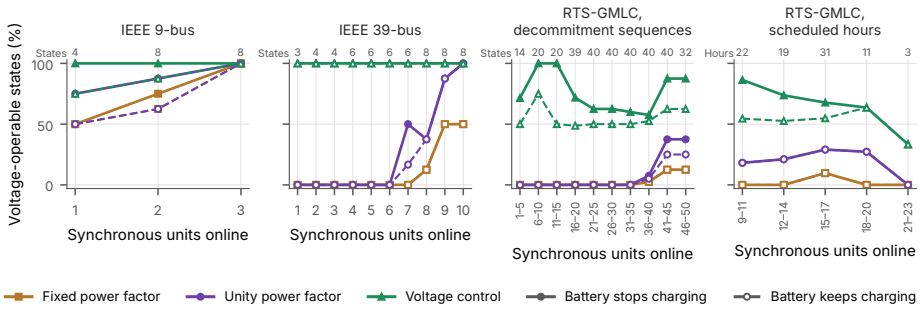


Fig. 8 Share of voltage-operable states against the synchronous units online, with the optimal setting, by plant mode. In the RTS-GMLC, units are grouped in fives in the sequences and in threes in the scheduled hours; numbers above points, states or hours per group. Dashed lines: battery that keeps charging

meets the ceiling before it. In the other three there is no admissible set point: in one the lowest voltage after the outage stays below 380 kV even at 420 kV; in another an outage has no solution at any set point; in the last the pre-event state has no solution below about 410 kV and exceeds the ceiling above it. With plants in voltage control, the lowest voltage after the outage rises in all four states, every outage has a solution, and three states have an admissible interval 10 to 22.5 kV wide.

The advantage of voltage control remains with a lower floor and with other droop slopes. A floor of 0.90 pu makes many more states operable at fixed and unity power factor (Fig. 10). Many of them are states that the plants pull down before any outage, which voltage control had already made operable. Voltage control still leads: in the RTS-GMLC sequences, 79% of the states are operable, against 31% at unity power factor. A droop slope of 2 or 7% instead of 3.33% changes the operable states far less than the control mode does. A smaller slope gives more reactive power for the same voltage change after the outage, and widens the interval of admissible set points

Table 3 Voltage-operable states with the optimal setting, by plant mode, reactive capability and battery behaviour. Const. Q: constant reactive power; VC: voltage control

		Battery stops charging						Battery keeps charging		
				Const.	VC					
Set of states	States	Fixed	Unity	Q	15%	30%	45%	Fixed	Unity	VC
9-bus, sequences	20	16	18	20	20	20	20	15	15	18
39-bus, sequences	59	9	21	54	59	59	59	9	19	59
RTS-GMLC, sequences	325	10	30	183	211	238	250	10	20	177
RTS-GMLC, scheduled hours	86	3	20	54	58	62	63	3	20	47

(Fig. 11). For an operator, that width is the margin left to choose the voltage schedule: a slope of 2% widens it by 2.5 kV, and one of 7% narrows it by up to 7.5 kV. The slope matters more in the 39-bus system, but with every slope voltage control keeps more states operable than unity power factor.

4.5 Co-located storage

A battery that keeps charging after the plant at its bus trips adds to the loss of injection in that area. It never made a state operable, and it removed operability from 19% of the combinations operable with the battery that stops, most of them in voltage control (Fig. 7, Table 3). In voltage control it removed 26% of the operable states in the RTS-GMLC sequences, 24% in the scheduled hours and none in the 39-bus system. In the RTS-GMLC sequences the losses occur where the batteries charge most: at the hour of largest resource with high renewable capacity, E2 and E4, the batteries charge at almost their full rating, a median of 3.8 GW before the outage. The sequences of least demand, with 1.6 GW, lose none (right panel of Fig. 4). This answers Q4.

To find which battery matters, each battery was left charging in turn while the others stopped. In every state that lost operability, a single battery was enough. In the 9-bus and 39-bus systems and in the scheduled hours, it is the battery at the bus of the outage with the largest violation, and stopping that battery alone restores operability. In the RTS-GMLC sequences, the batteries at buses 208 and 308 are each enough on their own. Both charge heavily and each feeds a radial load bus, 207 or 307. The outage of their own plant decides operability even when another outage has the largest violation. Stopping only the battery at the bus of that other outage restores operability in just 4% of the states that lost it and keep a load-flow solution. The batteries at buses 108 and 119 rarely mattered.

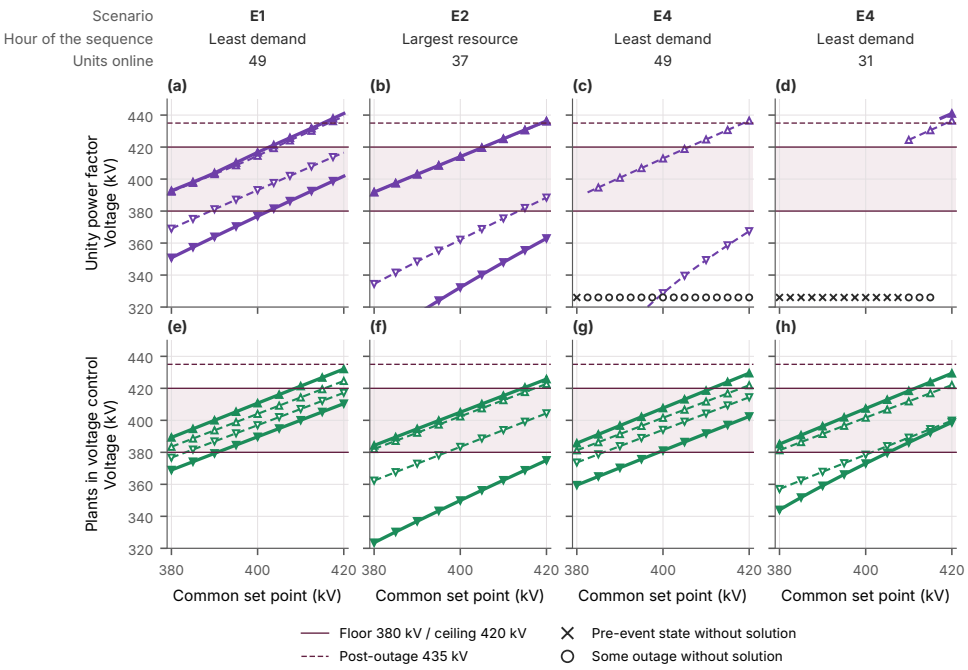


Fig. 9 Network voltages against the common set point in four states of the RTS-GMLC, at unity power factor (top) and in voltage control (bottom). Solid lines: after the worst outage; dashed lines: before it. Horizontal lines: 380, 420 and 435 kV. Crosses and circles: pre-event state or some outage without solution

4.6 Summary and Online Resource

The answers to Q1 to Q4, with the figures that show them and the tables of Online Resource 1 that give the counts, are:

- Q1, why a state fails. Almost every failed state lacks reactive power, whatever limit carries its violation: a capacitor reduces the violation in 99% of the combinations probed. Conflicts between the floor and the upper limits are 63% of those diagnosed. Fig. 3a shows it, and Tables S2 and S3 give the counts.
- Q1, upper limits. Where the largest violation lies on the upper limits, a capacitor still reduces it in 98% of the combinations probed. In the scheduled hours, a cable end rises above the ceiling because an area without units is raised to support a radial load bus. Figs. 5 and 6 show it, and Tables S2 and S9 give the counts.
- Q2, before or after an outage. At fixed or unity power factor, 60% of the failed combinations fail before any outage, and in voltage control 15%. Fig. 3b shows it by class, and Table S3 gives the counts.
- Q3, plant voltage control. No operable state becomes inoperable. The operable states rise from 9% to 73% in the RTS-GMLC sequences and from 23% to 72% in the scheduled hours, mostly through the reactive capability. With half the capability,

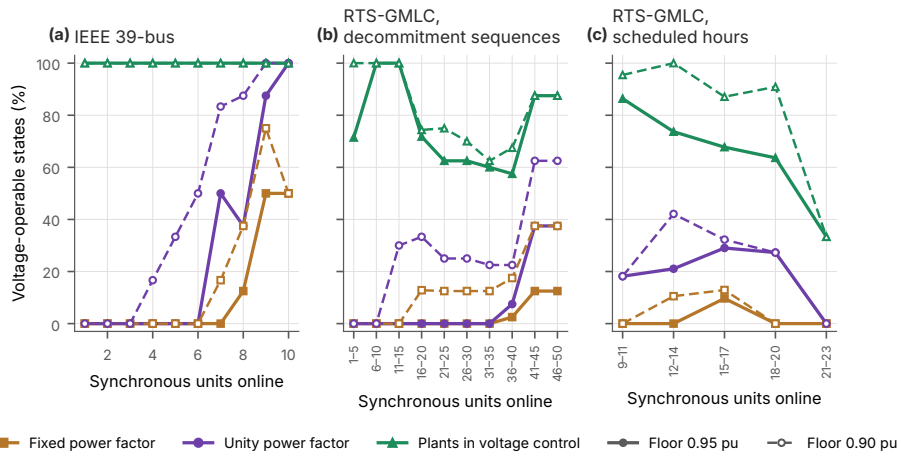


Fig. 10 Share of voltage-operable states against the units online, with the optimal setting and floors of 0.95 pu (solid) and 0.90 pu (dashed): (a) 39-bus system; (b) RTS-GMLC decommitment sequences; (c) RTS-GMLC scheduled hours

voltage control does at least as well as constant reactive power. The lead holds with a floor of 0.90 pu and with droop slopes of 2 and 7%. Figs. 7 to 11 and Table 3 show it, and Tables S1, S4, S5 and S8 give the counts.

- Q4, co-located battery. A battery that keeps charging removes 26% of the states operable in voltage control in the RTS-GMLC sequences and 24% in the scheduled hours. One battery is enough, and in the RTS-GMLC sequences it is usually not the one at the bus of the outage with the largest violation. Figs. 4 and 7 show it, and Tables S6 and S7 give the counts.

Table S1 gives the operable states per sequence, also with plants at constant reactive power. Tables S2 and S3 give the location of the largest violation and the diagnostic per set of states and plant mode. Tables S4 and S5 give the operable states with a floor of 0.90 pu and with droop slopes of 2 and 7%. Tables S6 and S7 give the effect of the battery that keeps charging, per set of states and plant mode and with one battery at a time. Table S8 gives the states that only voltage control makes operable, compared with constant reactive power, and Table S9 the scheduled hours per commitment rule.

5 Discussion

The main practical consequence is that, at low net load, the largest violation of a failed security check was not a reliable guide to the remedy. In the systems studied, the failures were mostly a lack of reactive power where it was needed. This held in the three systems, in the scheduled hours, and with plants in voltage control, which do supply reactive power. In long, lightly loaded 400 kV networks, such as those of Greece or the Iberian Peninsula, line charging can make an excess of reactive power

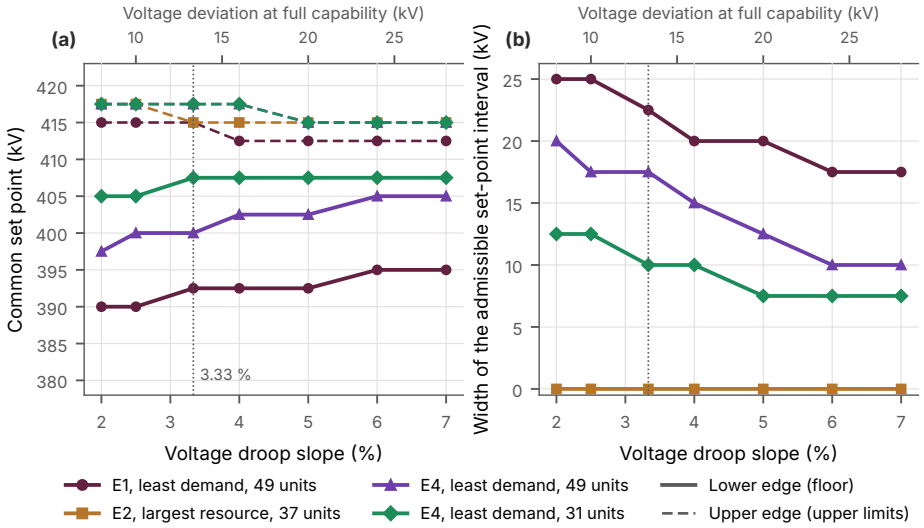


Fig. 11 (a) Lowest and highest admissible common set point and (b) width of the admissible interval against the droop slope of the plants in voltage control, in the four states of Fig. 9. Dotted line: 3.33%. Top axis: voltage deviation at full capability

the binding problem [2, 16]. The diagnostic tells the two situations apart in a given network.

The upper limits do bind in some states, but where the largest violation lay on them it was almost always a symptom. The best setting had to raise a whole area to keep the voltage up in a weak part of it. In the scheduled hours of the RTS-GMLC, the bus that crossed the ceiling, a cable end, was not the bus that failed. This is the local reactive deficiency that infeasibility analysis finds in collapsed load pockets [37], areas that network limits isolate from the rest of the system [39]. It also matches the operating experience that voltage violations were removed by committing nearby units [28]. An operator, or an optimisation, that reads the largest violation at face value may apply the opposite remedy: reactors at the cable end did not relieve any remedy hour and made the violation worse. The diagnostic costs about five more solutions of the same problem per state and gave a clear direction in almost every case.

Most failures at fixed or unity power factor came before any outage, and this matters for how security is checked at low net load. Contingency screening and a preventive SCOPF assume that the pre-outage state has an admissible voltage setting. At low net load, with few units online and plants that do not control voltage, that setting may not exist. The operating point itself then fails, not its response to an outage. Plant voltage control removed most of these failures, which is one more reason to require it.

Where support was missing depended on the commitment. In the E1 sequence of largest resource, the units at one radial load bus decided operability. In the scheduled hours, requiring at least eight large units online lowered the share of operable hours in voltage control. Planning studies reach the same conclusion from the other side:

the location of reactive sources shapes what can be built and retired [33–35]. For operation, the results point to remedies for areas left without a synchronous unit: plants in voltage control, synchronous condensers, retired units kept as condensers, and switched capacitors.

Of those remedies, this paper assesses plant voltage control. Most of its effect came from the reactive capability itself. Even so, plants in voltage control with half the capability did at least as well as plants holding the full capability at an optimised constant value. A requirement of capability alone, without voltage control, leaves part of the benefit unused. This agrees with the Nordic study of Kuhrmann et al. [9] and illustrates in numbers the recommendation of the Iberian investigation [1]. Grid codes differ widely in the capability they require, and in whether they require it at low production [40, 42]. A capability of 15% represents the weaker codes and still keeps most of the advantage.

Storage needs a requirement of its own. Guidance asks batteries to keep controlling voltage while charging [40], and local capacity studies treat their charging as load to be served under contingencies [43]. What matters here is what a battery does when the plant beside it trips. If it keeps charging, the outage of the plant becomes a larger loss of injection in an area that may already lack support. In the RTS-GMLC sequences, the battery that removed operability was usually not the one at the bus of the outage with the largest violation. A requirement set only for that outage is therefore not enough: every battery that shares a bus with a plant should stop charging when its plant trips.

The results hold for the assessment of Section 3.4 and for test systems. Three extensions would test them further. Dynamic simulation with relay timing and tap changers would check the states found operable within the protection window. Branch outages, especially of cables and long lines, would add the events most likely to cause an excess of reactive power. Corrective actions after the outage would enlarge the set of operable states.

6 Conclusions

This paper offers two tools for the security assessment of low-net-load states. The first, an operability test, poses voltage operability under single generation outages as the minimum-violation problem of a preventive SCOPF. Any tool that solves that problem can decide whether a state has a valid setting of the voltage controls, and load flows at frozen controls certify each positive verdict. The second is a diagnostic procedure that needs about five more solutions of the same problem. It tells which limit carries the violation of a failed state and whether the state lacks reactive power or has too much, and therefore which remedy applies. It does not use the structure of the network, so it can be applied to other systems as they are.

Applied to three test systems and to hours scheduled by a unit commitment, the two tools showed that:

1. A violation of the upper limits was mostly a symptom of lack of reactive support. It appears when an area without nearby units must be raised to keep the voltage

up at a radial load bus. Absorbing reactive power, the remedy that the largest violation suggests, made those states worse.
2. With plants at fixed or unity power factor, 60% of the combinations that could not be operated failed before any outage. A security check at low net load cannot assume that the pre-outage state is admissible.
3. Plant voltage control made no operable state inoperable. Compared with unity power factor, it raised the share of operable states from 9% to 73% in the RTS-GMLC sequences and from 23% to 72% in the scheduled hours, mostly through its reactive capability. With half the capability it did at least as well as the full capability held constant.
4. A battery that kept charging after its plant tripped removed 26% of the operable states in voltage control in the RTS-GMLC sequences. The decisive battery was usually not the one at the bus of the outage with the largest violation, so a requirement must cover every co-located battery.

Applying the diagnostic to long 400 kV networks, where an excess of reactive power may bind, would show whether it identifies that situation and its remedy as well.

Statements and Declarations

Funding. The authors did not receive support from any organisation for the submitted work.

Competing interests. The authors have no competing interests to declare that are relevant to the content of this article.

Data availability. The code, the exported low-net-load states, the result tables of every test and the scripts that compute each number of this paper and the data behind each figure are openly available in Zenodo, https://doi.org/10.5281/zenodo.23800264.

Supplementary information. Online Resource 1 (PDF) contains Tables S1 to S9, summarised at the end of Section 4.

Use of generative AI. Generative assistants based on large language models were used to help write, debug and test the simulation and analysis code, to search and organise the literature, and to edit the English text. Every number and result figure comes from the load flows and optimisations described above, computed by the authors, who checked each proposal before accepting it and take full responsibility for the content.

Author contributions. N. Moreno-Alfonso: conceptualization, methodology, software, formal analysis, investigation, data curation, visualization, writing – original draft. A. Bachiller-Soler and J. C. Bravo-Rodríguez: supervision, writing – review and editing. C. Álvarez-Arroyo: validation, writing – review and editing.

Ethics approval. Not applicable.

References

[1] Grid incident in Spain and Portugal on 28 April 2025 – final report. Final report of the ICS Investigation Expert Panel, ENTSO-E (March 2026). https://eepublicdownloads.blob.core.windows.net/public-cdn-container/clean-documents/Publications/2025/iberian-blackout/Final%20Report%20on%20the%20Grid%20Incident%20in%20Spain%20and%20Portugal%20on%2028%20April%202025.pdf Accessed 2026-09-23

[2] Rouco, L., Echavarren, F.M., Lobato, E.: The overvoltage-driven blackout of the Iberian Peninsula on 28th April 2025. Sustainable Energy, Grids and Networks **45**, 102125 (2026) https://doi.org/10.1016/j.segan.2026.102125

[3] Bialek, J.W.: The 2025 iberian blackout: Will there be more? Current Sustainable/Renewable Energy Reports **13**(1), 15 (2026) https://doi.org/10.1007/s40518-026-00297-w

[4] Vournas, C., Van Cutsem, T.: A review of overvoltage instability mechanisms. Sustainable Energy, Grids and Networks **47**, 102475 (2026) https://doi.org/10.1016/j.segan.2026.102475

[5] Comisión Nacional de los Mercados y la Competencia: Resolución de 30 de junio de 2026, por la que se modifican los procedimientos de operación eléctricos 7.4 y 14.4 para fomentar la prestación dinámica del servicio de control de tensión. Boletín Oficial del Estado, no. 167, 10 July 2026, pp. 96648–96692, BOE-A-2026-15127 (2026). https://www.boe.es/boe/dias/2026/07/10/pdfs/BOE-A-2026-15127.pdf Accessed 2026-10-19

[6] European Commission: Commission Regulation (EU) 2016/631 of 14 April 2016 establishing a network code on requirements for grid connection of generators. Official Journal of the European Union, L 112, 27 April 2016, pp. 1–68 (2016). https://eur-lex.europa.eu/eli/reg/2016/631/oj

[7] WECC Renewable Energy Modeling Task Force: WECC Guide for Representation of Photovoltaic Systems in Large-Scale Load Flow Simulations. Western Electricity Coordinating Council, Salt Lake City (2010)

[8] European Commission: Commission Regulation (EU) 2017/1485 of 2 August 2017 establishing a guideline on electricity transmission system operation. Official Journal of the European Union, L 220, 25 August 2017, pp. 1–120 (2017). https://eur-lex.europa.eu/eli/reg/2017/1485/oj

[9] Kuhrmann, L., Chen, P., Göransson, L., Johnsson, F.: Assessment of the N-1 voltage security of a future Nordic energy system. arXiv:2609.24356 (2026). https://doi.org/10.48550/arXiv.2609.24356

[10] Prionistis, G., Vournas, C.: Flexibility region discontinuity due to overvoltage

constraints. Sustainable Energy, Grids and Networks **43**, 101807 (2025) https://doi.org/10.1016/j.segan.2025.101807

[11] Monticelli, A., Pereira, M.V.F., Granville, S.: Security-constrained optimal power flow with post-contingency corrective rescheduling. IEEE Transactions on Power Systems **2**(1), 175–180 (1987) https://doi.org/10.1109/TPWRS.1987.4335095

[12] Capitanescu, F., Martinez Ramos, J.L., Panciatici, P., Kirschen, D., Marano Marcolini, A., Platbrood, L., Wehenkel, L.: State-of-the-art, challenges, and future trends in security constrained optimal power flow. Electric Power Systems Research **81**, 1731–1741 (2011) https://doi.org/10.1016/j.epsr.2011.04.003

[13] Castillo, A., Laird, C., Silva-Monroy, C.A., Watson, J.-P., O'Neill, R.P.: The unit commitment problem with AC optimal power flow constraints. IEEE Transactions on Power Systems **31**(6), 4853–4866 (2016) https://doi.org/10.1109/TPWRS.2015.2511010

[14] Chu, Z., Teng, F.: Voltage stability constrained unit commitment in power systems with high penetration of inverter-based generators. IEEE Transactions on Power Systems **38**(2), 1572–1582 (2023) https://doi.org/10.1109/TPWRS.2022.3179563

[15] Davoodi, E., Capitanescu, F., Wehenkel, L.: A methodology to evaluate reactive power reserves scarcity during the energy transition. IEEE Transactions on Power Systems **38**(5), 4355–4368 (2023) https://doi.org/10.1109/TPWRS.2022.3216639

[16] Mandoulidis, P., Lampropoulos, V., Vournas, C., Karystianos, M., Christoforidis, G., Neris, A., Kabouris, Y.: Methods to mitigate voltage-rise issues in transmission systems. Electric Power Systems Research **236**, 110888 (2024) https://doi.org/10.1016/j.epsr.2024.110888

[17] Albustami, A.A., Taha, A.F.: The Iberian Blackout: A Black Swan or a Gray Rhino? A Protection-Aware Dynamic Voltage Security Assessment. arXiv:2511.17433 (2026). https://doi.org/10.48550/arXiv.2511.17433

[18] Vournas, C., Van Cutsem, T.: On overvoltage instability. Electric Power Systems Research **262**, 113589 (2027) https://doi.org/10.1016/j.epsr.2026.113589

[19] Morão, H.: The 2025 Iberian Peninsula blackout: Lessons for modern power systems and policy implications. Utilities Policy **101**, 102196 (2026) https://doi.org/10.1016/j.jup.2026.102196

[20] Anderson, P.M., Fouad, A.A.: Power System Control and Stability. 2nd edn. Wiley-IEEE Press, Piscataway, NJ (2003). https://doi.org/10.1109/9780470545577

[21] Athay, T., Podmore, R., Virmani, S.: A Practical Method for the Direct Analysis of Transient Stability. IEEE Transactions on Power Apparatus and Systems **PAS-98**(2), 573–584 (1979) https://doi.org/10.1109/TPAS.1979.319407

[22] Barrows, C., Preston, E., Staid, A., Stephen, G., Watson, J.-P., Bloom, A., Ehlen, A., Ikaheimo, J.: The IEEE reliability test system: A proposed 2019 update. IEEE Transactions on Power Systems **35**(1), 119–127 (2020) https://doi.org/10.1109/TPWRS.2019.2925557

[23] Daly, P., Flynn, D., Cunniffe, N.: Inertia considerations within unit commitment and economic dispatch for systems with high non-synchronous penetrations. In: 2015 IEEE Eindhoven PowerTech, pp. 1–6 (2015). https://doi.org/10.1109/PTC.2015.7232567

[24] Vournas, C.D., Nikolaidis, V.C., Tsourakis, G.I.: Coordinated countermeasures against overvoltage instability in autonomous power systems. IEEE Transactions on Power Delivery **36**(6), 3329–3338 (2021) https://doi.org/10.1109/TPWRD.2020.3038773

[25] Van Cutsem, T.: Voltage instability: phenomena, countermeasures, and analysis methods. Proceedings of the IEEE **88**(2), 208–227 (2000) https://doi.org/10.1109/5.823999

[26] Wächter, A., Biegler, L.T.: On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Mathematical Programming **106**(1), 25–57 (2006) https://doi.org/10.1007/s10107-004-0559-y

[27] Andersson, J.A.E., Gillis, J., Horn, G., Rawlings, J.B., Diehl, M.: CasADi: a software framework for nonlinear optimization and optimal control. Mathematical Programming Computation **11**(1), 1–36 (2019) https://doi.org/10.1007/s12532-018-0139-4

[28] Lobato, E., Rouco, L., Gómez, T., Echavarren, F.M.: Elimination of voltage violations in the Spanish electricity market. Part I: Theory. Electric Power Systems Research **75**(2–3), 99–109 (2005) https://doi.org/10.1016/j.epsr.2004.12.006

[29] Lobato Miguélez, E., Rouco Rodríguez, L., Gómez San Román, T., Echavarren Cerezo, F.M., Navarrete Fernández, M.I., Casanova Lafarga, R., López Camino, G.: A practical approach to solve power system constraints with application to the Spanish electricity market. IEEE Transactions on Power Systems **19**(4), 2029–2037 (2004) https://doi.org/10.1109/TPWRS.2004.831704

[30] Corsi, S., Pozzi, M., Sabelli, C., Serrani, A.: The coordinated automatic voltage control of the Italian transmission grid—Part I: Reasons of the choice and overview of the consolidated hierarchical system. IEEE Transactions on Power Systems **19**(4), 1723–1732 (2004) https://doi.org/10.1109/TPWRS.2004.836185

[31] Paul, J.P., Leost, J.Y., Tesseron, J.M.: Survey of the secondary voltage control in France: Present realization and investigations. IEEE Transactions on Power Systems **2**(2), 505–511 (1987)

[32] Omi, S., Shirai, Y.: Benders-decomposition-based voltage-stability-constrained unit-commitment taking reactive power limit into account. IEEJ Transactions on Electrical and Electronic Engineering **18**, 375–385 (2023) https://doi.org/10.1002/tee.23733

[33] Wogrin, S., Tejada-Arango, D., Delikaraoglou, S., Botterud, A.: Assessing the impact of inertia and reactive power constraints in generation expansion planning. Applied Energy **280**, 115925 (2020) https://doi.org/10.1016/j.apenergy.2020.115925

[34] Wen, Y., Yan, G., Deng, B., Xing, P., Yu, L., Jiang, X.: Coordination strategy of coal-fired units decommissioning and condenser configuration considering the support capability constraints of power systems. Advanced Engineering Sciences **55**(1), 26–34 (2023). In Chinese https://doi.org/10.15961/j.jsuese.202200613

[35] Masood, N.-A., Yan, R., Saha, T.K., Bartlett, S.: Post-retirement utilisation of synchronous generators to enhance security performances in a wind dominated power system. IET Generation, Transmission & Distribution **10**(13), 3314–3321 (2016) https://doi.org/10.1049/iet-gtd.2016.0267

[36] Australian Energy Market Operator: Transfer limit advice – Minimum generator combinations in SA. Version 46. AEMO, Melbourne (2025)

[37] Jereminov, M., Bromberg, D.M., Pandey, A., Wagner, M.R., Pileggi, L.: Evaluating feasibility within power flow. IEEE Transactions on Smart Grid **11**(4), 3522–3534 (2020) https://doi.org/10.1109/TSG.2020.2966930

[38] Platbrood, L., Capitanescu, F., Merckx, C., Crisciu, H., Wehenkel, L.: A generic approach for solving nonlinear-discrete security-constrained optimal power flow problems in large-scale systems. IEEE Transactions on Power Systems **29**(3), 1194–1203 (2014) https://doi.org/10.1109/TPWRS.2013.2289990

[39] Lesieutre, B.C., Thomas, R.J., Mount, T.D.: Identification of load pockets and market power in electric power systems. Decision Support Systems **40**, 517–528 (2005) https://doi.org/10.1016/j.dss.2004.09.010

[40] North American Electric Reliability Corporation: Reliability Guideline: BPS-Connected Inverter-Based Resource Performance. NERC, Atlanta (2018)

[41] North American Electric Reliability Corporation: Transmission System Planning Performance Requirements, Reliability Standard TPL-001-5.1. NERC, Atlanta (2020)

[42] Sarkar, M.N.I., Meegahapola, L.G., Datta, M.: Reactive power management in renewable rich power grids: A review of grid-codes, renewable generators, support devices, control strategies and optimization algorithms. IEEE Access **6**, 41458–41489 (2018) https://doi.org/10.1109/ACCESS.2018.2838563

[43] California ISO: 2031 Local Capacity Technical Study – Final Report and Study Results. California ISO, Folsom (2026)

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- eesupplementary.pdf