THE IBERIAN BLACKOUT: A Black Swan OR A Gray Rhino?† A THOROUGH POWER SYSTEM ANALYSIS

Abdallah Alalem Albustami and Ahmad F. Taha*

Abstract—On April 28, 2025, the Iberian power system suffered a full blackout. It was the first documented overvoltage-driven cascade in Europe. The event sparked debate about root causes, including high renewables output, low inertia, and operator actions. This paper presents a thorough power system analysis of the incident to sort signal from noise and explain, step by step, how the blackout unfolded. Specifically, we (i) reconstruct the timeline and causal chain of the incident, (ii) present and summarize contributing factors using factual findings from incident reports, (iii) reproduce the blackout on an IEEE test system, (iv) analyze the incident from a system-theoretic, voltage-control perspective, and (v) translate our analysis into practical, technical measures that aim to mitigate and prevent similar incidents.

Index Terms—Voltage stability, power system protection, blackout analysis, renewable integration.

I. INTRODUCTION AND PAPER CONTRIBUTIONS

On April 28, 2025, at 12:33:27 CEST, the Iberian mainland power system experienced a complete blackout, disconnecting Spain, Portugal, and a small region of France near the Spanish border. The cascade from initial overvoltage-induced generation trips to system-wide collapse occurred within approximately 30 seconds.

This blackout differs fundamentally from historical patterns. Major grid failures typically originate from either extreme natural events or from transmission system overload followed by undervoltage or underfrequency cascades. It represents the first documented case in Continental Europe where cascading generation disconnections occurred due to sustained overvoltage conditions at collector substations that remained unobservable to transmission operators, a failure mode unprecedented in the synchronous area. The incident prompted considerable discussion and debate regarding root causes, with analyses variously emphasizing high renewable generation, system inertia, voltage control inadequacies, and operational decisions made to address morning oscillations. Rather than advocate for a single causal narrative, this work seeks to present a thorough, system-theoretic understanding of the different blackout causes.

Paper Audience, Contributions, and Organization. The only required background needed for the readability of this paper are basic power system analysis and some rudimentary control systems language. The paper contributions are as follows.

 While many incident reports describe the incident's chronology in detail, this work adds a complementary perspective by analyzing how the documented factors interacted to produce the cascade by identifying the enabling mechanisms and their dynamic coupling via a differential-algebraic power system model.

- The cascade is replicated on the IEEE 39-bus system to enable diagnostics and a better system-level understanding of how the blackout unfolded. All simulation codes are publicly available [1].
- The cascade is analyzed through the differential-algebraic structure governing transmission systems to reveal how topology changes and inadequate reactive response degraded voltage control on the timescales relevant for protection coordination.
- Three practical measures are proposed—finite-horizon voltage control assessment, stability-aware meshing guidance, and voltage-aware emergency coordination. These could potentially mitigate similar over-voltage cascades.

The remainder of this paper is organized as follows. Section III reconstructs the Iberian blackout chronology. Section III describes how we replicate the cascade on a test system. Section IV discusses the factors that led to the blackout using factual findings from incident reports and our understanding. Section V develops a system-theoretic, voltage control analysis of the blackout, and Section VI presents replication results. Section VII proposes practical mitigation measures and concludes the paper.

II. THE IBERIAN BLACKOUT: EVENT SEQUENCE

This section delineates the event sequence that led to the blackout and summarizes the main factual findings documented in official reports in technical detail, for a high-level overview, refer to Fig. 1. The European Network of Transmission System Operators for Electricity (ENTSO-E) expert panel published a 262-page factual report [2] providing a detailed sequence of events. Red Eléctrica de España (REE) also released a technical report on the incident [3], and the Spanish government published an analysis committee report in Spanish [4]. These documents establish the event chronology but explicitly defer root-cause analysis and recommendations to future work. We refer to these as the primary sources.

The morning of April 28 was characterized by light system loading and high renewable integration. At 12:30, the generation mix comprised 82% renewables, predominantly solar and wind inverters, 10% nuclear, 3% combined-cycle gas, 1% coal, and 4% cogeneration, meeting light holiday demand of 25,184 MW (around 56% of historical peak) [4, p. 38]. System inertia was correspondingly low at $H_{\rm tot}$ =2.21–2.71 seconds [2, p. 36].

Minor inter-area oscillations (0.2 Hz, causing 4–7 kV voltage swings) occurred at 10:30, 11:03, and 11:23, all of which damped naturally within minutes [3, p. 3]. Between 12:03 and 12:08, a forced oscillation with dominant frequency 0.63 Hz emerged in southwestern Iberia, producing voltage swings reaching 30 kV at some 400 kV substations and exciting the 0.21 Hz East-Centre-West inter-area mode [2, p. 53]. The origin of the 0.63 Hz oscillation remains unresolved (at the time of writing in November 2025), REE characterizes it as a *forced* oscillation possibly originating at

[†]A *Black Swan* is a rare, hard-to-predict event with outsized impact that seems *obvious* only in hindsight. In contrast, a *Gray Rhino* is a highly probable, high-impact threat that is visible and approaching, yet neglected until it happens.

^{*}Corresponding author. A. Alalem Albustami and A. F. Taha are researchers at Vanderbilt University, Nashville, TN 37235, USA. (e-mails: {abdallah.b.alalem.albustami, ahmad.taha}@vanderbilt.edu).

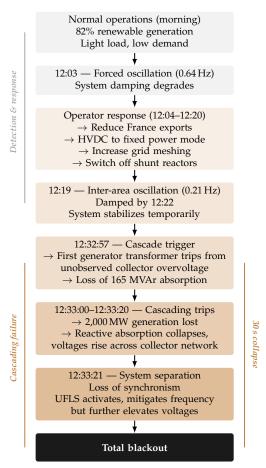


Figure 1: High-level chronology of the April 28 blackout.

a PV plant in Badajoz [3, p. 12], whereas the ENTSO-E Expert Panel did *not* confirm the cause [2, p. 59]. A second inter-area oscillation (0.21 Hz) occurred between 12:19 and 12:22, with the Iberian Peninsula oscillating coherently against Continental Europe.

Operators responded per protocols: between roughly 12:03 and 12:25, they energized (energizing a circuit, also referred to as network meshing, is to close breakers and apply voltage so the line is in service) several 400 kV circuits to decrease system impedance, switched off several shunt reactors during the oscillation window, with the first reconnections occurring minutes later (at around 12:26) as voltages began to rise [2, pp. 19–21]. At 12:05, exports to France were reduced, and at 12:19 REE asked RTE to reduce the exchange further. At 12:11, operators reconfigured the HVDC interconnection from AC-emulation (PMODE3) to fixed-power mode (PMODE1*) at 1,000 MW export [4, p. 103]. These actions damped oscillations by 12:22, and by 12:30 voltages settled to 410-420 kV (within the 375-435 kV envelope) and frequency stabilized near 50 Hz. At 12:26, recognizing deteriorating conditions, operators requested rapid synchronization of additional conventional generators with dynamic voltage control capability in the southern zone. The shortest offered lead time was 1.5 hours, meaning the unit could not be online until around 14:00—far beyond the sub-minute timescale of the impending cascade [4, p. 35].

*PMODE1 (fixed-power): line controls active power to a constant setpoint regardless of system frequency/angle. PMODE3 (AC-emulation): active power tracks voltage-angle difference so the HVDC contributes inter-area damping.

Between 12:32:00 and 12:32:57, approximately several hundred megawatts were lost from distributed generation and net load increases. At 12:32:57, a 400/220 kV generation transformer tripped due to overvoltage on the 220 kV collector side (242 kV), disconnecting 355 MW and 165 MVAr of reactive absorption transmission voltage was 418 kV (within limits), but the tap position had not adjusted as voltages evolved, creating secondary-side overvoltage invisible to transmission measurements [2, p. 28]. This initiated a self-reinforcing cascade: reactive absorption loss elevated voltages, reduced power flows decreased line reactive consumption, and P.O. 7.4[†] generators exhibited threshold-type response rather than continuous absorption [2, p. 76]. Within 20 seconds, approximately 2,000 MW disconnected. The pattern was consistent across events: transmission voltages stayed within or marginally above limits, but internal plant voltages downstream of transformers with lagged tap positions exceeded local protection settings.

The massive generation loss created acute northern-region deficit. France imports surged to 3.8 GW by 12:33:19, sudden power reversal and angle stress caused loss of synchronism at 48.46 Hz, and AC interconnections tripped correctly at 12:33:21, islanding Iberia and preventing Continental Europe propagation [2, pp. 108–109]. The islanded system faced dual instability: southern overvoltage and system-wide frequency collapse. Underfrequency load shedding arrested the frequency fall but, by removing reactive consumption, pushed voltages higher and accelerated the overvoltage cascade [3, p. 11]. At 12:33:27, transmission voltages collapsed and the system was *de-energized*, marking a complete blackout.

III. BLACKOUT REPLICATION

To systematically investigate the cascade mechanisms identified in incident reports, we replicate the April 28 event on the IEEE 39-bus New England test system using ANDES, an open-source Python software for symbolic power-system modeling and numerical analysis [5]. The replication enables controlled examination of individual cascade drivers and serves two purposes: (i) validating that the identified failure modes can produce cascading overvoltage collapse under conditions matching the incident, and (ii) providing a reproducible framework for the analyses presented throughout this paper. Complete simulation framework including all intricate details, ANDES case files and figure generation scripts, is available on GitHub [1].

Replicating a continental-scale blackout on a test system presents several technical challenges, which we address as follows: (i) Collector-level voltage decoupling: The IEEE 39-bus system contains no explicit collector substations or generation step-up transformers with time-lagged tap positions, however incident reports identify the transmission-distribution voltage decoupling as critical to the cascade. We address this by augmenting the base case with explicit collector buses at 138 kV connected to transmission buses (345 kV) through transformers with fixed tap ratios. These fixed taps represent on-load tap changer positions set during morning voltage fluctuations that did not adjust as transmission voltages subsequently rose. When transmission voltages increase due to operator actions, the lagged tap positions amplify collector-side voltages beyond protection thresholds, replicating the mechanism where transmission measurements remained within

 $^\dagger PO.$ 7.4: operating procedure that governs the ancillary voltage-control service. It obliges designated resources to regulate voltages continuously at TSO-assigned control nodes through measures such as automatic voltage regulators (AVRs).

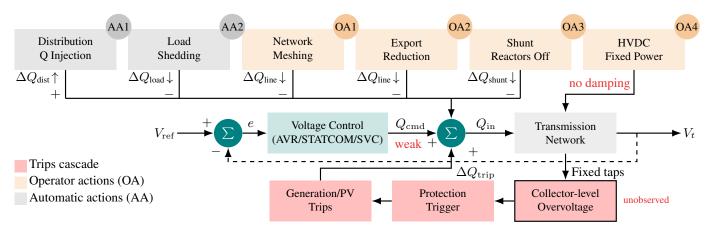


Figure 2: Feedback structure of the cascading overvoltage collapse. The negative feedback loop (dashed) through voltage control failed to stabilize due to limited AVR response, while the positive feedback loop (solid) through collector-level overvoltage and protection-triggered trips reinforced voltage rise. Operator actions (orange) inadvertently amplified the cascade by reducing natural reactive absorption.

operational limits while collector-side voltages exceeded trip settings. (ii) Aggregate generation representation: The Iberian system comprised over 20,000 MW of distributed generation across hundreds of substations, each with local protection and reactive devices. Explicitly modeling this granularity is computationally intractable and would require plant-specific parameters unavailable in public reports. We therefore create five representative collector groups, each equipped with inverter-based generation blocks and shunt reactors providing 150-300 MVAr of reactive absorption. When collector protection operates, the associated generation and reactive devices disconnect, and ANDES computes the resulting voltage response through standard power-flow equations. This captures the essential physics, each trip removes reactive absorption, the algebraic network constraints force voltage increases throughout the system, and elevated voltages trigger subsequent collector trips in a self-reinforcing cascade. (iii) Device rating calibration: Device ratings in the base IEEE 39-bus case produce lower voltage sensitivities than observed in the actual incident, where operator actions had measurable but not immediately alarming effects on transmission voltages. To ensure cascade drivers produce observable transmission-level responses matching incident phenomenology while remaining physically plausible, we scale certain device ratings (shunt reactors, parallel line reactances) to reflect aggregate regional effects documented in incident reports.

The simulation initializes the IEEE 39-bus case with high renewable penetration (80% of generation from inverter-based sources), reduced loads matching light holiday conditions (58% of nominal), and collector protection thresholds calibrated to 1.04–1.10 pu on the 138 kV base. Scheduled operator actions execute per incident chronology, collector voltages are monitored continuously, and when thresholds are exceeded for the configured dwell time, protection logic disconnects the associated generation and reactive devices. All voltage responses are simulated by ANDES through solving the network differential-algebraic model—presented later in (1).

All figures presented throughout this paper are generated from this replication framework with parameter variations to isolate specific effects. The replication intentionally simplifies several aspects: inter-area oscillations that preceded the cascade are not modeled, and protection coordination details are representative rather than plant-specific due to data unavailability. These simplifications enable computational tractability while preserving the essential

cascade structure identified in incident reports: transmission-collector voltage decoupling through lagged taps, reactive power imbalance propagation through the algebraic network, and sequential protection-triggered trips. We welcome any critique of essential components that we have not considered in the simulations; the Github codes are an initial step for further analyses [1].

IV. WHY IT HAPPENED:

ROOT-CAUSE IDENTIFICATION FROM INCIDENT REPORTS

From the three reports [2]–[4] we identified three key factors that created conditions for the cascading voltage collapse, as discussed next. The reports are explicitly factual and document the event sequence without root-cause analysis, which the expert panels deferred to future work [2, p. 2]. This contribution of this section is a synthesis of the findings to identify the blackout-enabling mechanisms. In this section and what follows, we refer to *operator actions* by (OA) and to *automatic actions* by (AA). Fig. 2 delineates the main OA/AA elements, shows how they interact, and highlights the overall feedback structure and cascade mechanism.

Unintended consequences of protocols. The measures taken to dampen the 0.63 Hz and 0.21 Hz oscillations, while procedurally correct, altered system structure in ways operators could not assess. Operators energized parallel 400 kV circuits to reduce system impedance and improve damping but this simultaneously reduced natural line reactive consumption, raising voltages. This is also referred to as network meshing which altered the topology of the network, we refer to these meshing actions as OA1 throughout. Reduced exports to France (1,300 MW) and Portugal (545 MW) which decreased transmission line loading [3, p. 5], further reducing reactive losses and elevating voltages, referred to as OA2. Additionally, operators disconnected shunt reactors throughout the morning to counteract transient undervoltages during oscillations, removing critical tools for subsequent voltage control (OA3) [3, p. 4]. The HVDC link was reconfigured from AC-emulation mode to fixed 1,000 MW export eliminated its ability to respond to frequency deviations: when the cascade began and frequency dropped, the link maintained constant export rather than reducing power to support the Iberian system (OA4).

Collector-level observability gap. Transmission voltages generally remained within operational limits (375–435 kV), yet generation facilities tripped en masse. Many plants tripped below the

440 kV/60-minute threshold or lacked adequate protection margins and time delays. What happened is that transformer tap positions had not adjusted as voltage profiles changed during the morning oscillations. When transmission voltage was 418 kV at 12:32:57, the 220 kV collector side measured 242 kV exceeding the transformer protection setting. Similar mismatches between transmission measurements and internal plant voltages, hidden behind lagged taps, drove subsequent trips. ENTSO-E notes that *the grid is not fully observable with high-resolution monitoring systems* [2, pp. 24, 115]. In practice, this meant the system satisfied transmission-level constraints while numerous collector substations violated their protection thresholds behind unmonitored transformers, creating a cascade mechanism invisible to operators until generation loss became widespread. We refer this as *the mismatch in observability*.

Voltage control failure. Generators mandated under P.O. 7.4 to provide continuous dynamic voltage control operated with a dead-band response rather than smooth droop-based regulation. Specifically, when voltage was between 405–410 kV (the regulatory setpoint band), generators provided no reactive response. REE's report explicitly states: Generation subject to Operating P.O. 7.4 failed to comply with its dynamic voltage control obligations... Generators typically respond only when voltage deviations become significant, suggesting that their response is primarily driven by internal plant protection mechanisms [3, p. 13]. The reason generators were obligated to act only outside the 405–410 kV band is likely (i) to avoid controller hunting and VAR loop flows from small deviations/measurement noise—thereby reducing losses and exciter wear—and (ii) to preserve reactive headroom and enable hierarchical coordination (secondary voltage control, shunts/STATCOMs) while keeping the service auditable via a clear performance band. In normal conditions this improves stability and operability, but in high-IBR, light-load situations can create a short-horizon dead zone unless complemented by fast absorbing devices. At 82% renewable penetration where inverters provide no dynamic voltage support, the mandated 18% of conventional generation was expected to shoulder the entire burden of continuous voltage regulation, which is a responsibility many units did not fulfill as transmission voltages remained mostly within that band. Several large generators in the southern zones reportedly absorbed insufficient reactive power or injected reactive power when absorption was required, contributing to further voltage rise.

Misattributed narratives: inertia and renewables. Public discourse following the blackout frequently emphasized low system inertia and high IBR penetration as root causes. While these factors shaped system dynamics, they misdiagnose the fundamental failure mode. The cascade initiated through overvoltage at 12:32:57 when system frequency remained near nominal—this was a voltage-driven event, not an inertia-limited frequency collapse. The subsequent frequency decline resulted from cascading generation losses triggered by overvoltage protection, not vice versa. High inertia would have slowed rate of change of frequency after trips began but could not have prevented the voltage-driven cascade initiation [3, p. 14]. Similarly, while 82% IBR penetration meant only 18% of conventional generation remained available for dynamic reactive control, this alone did not predetermine the outcome. Had the remaining conventional generators absorbed reactive power as required, and had operators possessed tools to assess how topology changes affected voltage stability margins, the cascade could have been prevented regardless of IBR

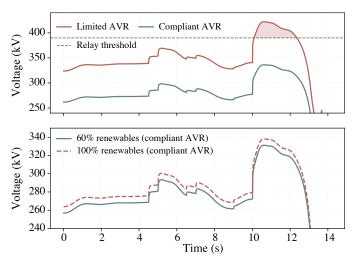


Figure 3: Effect of AVR compliance on overvoltage. Top: observed trajectory with limited AVR (red) versus generators providing continuous droop-based absorption (teal). Bottom: with compliant AVR, even higher renewable shares (60% and 100% lines) keep voltages below the relay limit across the same disturbance sequence.

penetration level. The critical failures were inadequate reactive control compliance, limited collector-level observability, and absence of dynamic voltage control assessment capability, all operational and infrastructural gaps and not inherent technological constraints of renewable integration. Fig. 3 illustrates this point directly: with continuous droop-based AVR, the peak stays below the relay limit, whereas limited AVR exceeds it under the same renewable share. With AVR compliance in place, even increasing renewables from 60% to 100% preserves headroom, underscoring that fast reactive control—not the renewable share itself—primarily governs overvoltage risk.

These reported facts establish three factors of the cascade: (i) a topology modification that reduced natural reactive consumption (meshing or energization), (ii) distribution-side overvoltages unobservable at the transmission layer, and (iii) weak seconds-scale reactive response. This does not explain the dynamic mechanism through which they interacted. The next section produces a thorough analysis on how those ingredients interacted dynamically from a system-theoretic perspective, complementing the empirical findings of this section.

V.

WHY IT HAPPENED: THE DYNAMIC VOLTAGE CONTROL GAP

The blackout can be examined from multiple perspectives. The analysis we present here focuses on one particular angle: how voltage and reactive power interact when network topology changes via meshing, and why operators lacked tools to assess whether their topology decisions would improve or worsen voltage control on the timescales relevant for preventing cascading trips.

A. Transmission Power System Model

Transmission systems are typically described by differential algebraic equations (DAEs) that couple plant dynamics with network constraints [6]. The dynamic state vector $\boldsymbol{x}_d(t) \in \mathbb{R}^{n_d}$ collects device-level physics, such as synchronous machine rotor angle, speed and transient EMFs together with AVRs, FACTS (SVC/STATCOM) internal states, and IBR states. The algebraic state vector $\boldsymbol{x}_a(t) \in \mathbb{R}^{n_a}$ comprises network variables that satisfy instantaneous power-balance constraints, typically bus voltage magnitudes,

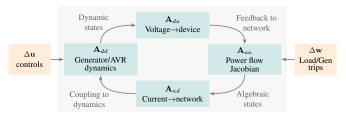


Figure 4: Block structure of the linearized DAE model in (2). Controls Δu excite device dynamics A_{dd} , devices and network are coupled by A_{da} and A_{ad} , the algebraic network is governed by the power-flow Jacobian A_{aa} , exogenous disturbances Δw act on the algebraic part.

angles, and active and reactive power flows. Control inputs $u(t) \in \mathbb{R}^{n_u}$ include field voltages, AVR reference voltages, and HVDC setpoint/mode. Exogenous inputs $w(t) \in \mathbb{R}^{n_w}$ include loads and renewables injections, voltage-dependent ZIP loads, net distribution-level injections, and protection-triggered P/Q changes. Functions $f(\cdot)$ and $g(\cdot)$ collect device-level differential equations and the algebraic AC network flow enforcing KCL/KVL. The model is given by:

generator dynamics:
$$\dot{x}_d(t) = f(x_d(t), x_a(t), u(t))$$
 (1a)
network power flow: $0 = g(x_d(t), x_a(t), u(t), w(t))$. (1b)

Linearizing around $(\bar{x}_d, \bar{x}_a, \bar{u}, \bar{w})$ yields a linearized DAE:

$$\underbrace{\begin{bmatrix} I & 0 \\ 0 & 0 \end{bmatrix}}_{E} \begin{bmatrix} \Delta \dot{x}_d \\ \Delta \dot{x}_a \end{bmatrix} = \underbrace{\begin{bmatrix} A_{dd} & A_{da} \\ A_{ad} & A_{aa} \end{bmatrix}}_{A} \begin{bmatrix} \Delta x_d \\ \Delta x_a \end{bmatrix} + \underbrace{\begin{bmatrix} B_d \\ B_a \end{bmatrix}}_{B_{uv}} \Delta u + \underbrace{\begin{bmatrix} 0 \\ B_w \end{bmatrix}}_{B_{uv}} \Delta w. \quad (2)$$

Here, A_{aa} is the power-flow Jacobian, A_{dd} captures device internal dynamics, and A_{ad} , A_{da} couple devices and network (how device currents affect nodes, and how nodal voltages/angles drive device controls). Input matrices B_u and B_w map control actions (field voltages, setpoints) and disturbances (load/generation trips) into the differential and algebraic subsystems, respectively. Fig. 4 visualizes this coupled structure.

To analyze the Iberian overvoltage cascade, we examine how the DAE model (2) responds to operator actions (meshing, switching off reactors, etc) and disturbances (generation trips). The analysis presented herein identifies the enabling conditions for instability at the *cascade's onset*, while the nonlinear replication in Section III and results in Section VI confirm that these conditions reproduce the time-domain behavior.

The objective of this analysis is to explain the mechanism by which operator actions degraded voltage control, not to predict the complete nonlinear cascade trajectory. The linearized DAE model (2), (i) reveals the feedback structure that enabled instability to initiate, which is the operationally relevant question for preventing similar events, and (ii) it quantifies the directional sensitivities and causal pathways through which reactive power disturbances affect voltages, giving operators explicit relationships between topology decisions and voltage control authority that remain valid for the small-to-moderate deviations at cascade onset, before protection actions drive the system into large-signal behavior.

We now reinterpret the chronology from a DAE-based voltagecontrol perspective, identifying how these structural and control changes interacted to enable the cascade.

B. Algebraic Imbalance and Dynamic Compensation Failure

The fundamental mechanism driving the overvoltage cascade lies in the coupling between algebraic power balance constraints and dynamic reactive response capability. The abstract algebraic model (1b) enforces instantaneous active and reactive power balance through:

$$P_{Gi} + P_{Ri} - P_{Li} = \sum_{j=1}^{N} v_i v_j \left(G_{ij} \cos \theta_{ij} + B_{ij} \sin \theta_{ij} \right), \qquad (3a)$$

$$Q_{Gi} + Q_{Ri} - Q_{Li} = \sum_{j=1}^{N} v_i v_j (G_{ij} \sin \theta_{ij} - B_{ij} \cos \theta_{ij}), \quad (3b)$$

where $v_i = |V_i|$ is the per–unit voltage magnitude at bus i, $\theta_{ij} = \theta_i - \theta_j$ is the voltage–angle difference, and $G_{ij} + jB_{ij}$ are conductance and susceptance entries of the bus admittance matrix \mathbf{Y} . P_{Gi} , Q_{Gi} denote synchronous generator injections at bus i, P_{Ri} , Q_{Ri} IBR injections, and P_{Li} , Q_{Li} net load consumptions (ZIP effects included).

When the 400/220 kV transformer tripped at 12:32:57, removing 355 MW and 165 MVAr of reactive absorption, equation (3b) could only remain satisfied through two mechanisms: (i) dynamic compensation via increased generator reactive output $\Delta Q_i^{\rm gen}>0$ through, for example, automatic voltage regulator (AVR) field current adjustments, or (ii) voltage rise throughout the network to rebalance the right-hand side summation.

The critical failure was the absence of mechanism (i): generators exhibited dead-band response rather than continuous proportional control. On the seconds-scale, dynamic reactive response is well summarized by a static droop with a short lag: outside the regulated setpoint band, deviations above $V_{\rm ref}$ command absorption and deviations below $V_{\rm ref}$ command injection; the conclusions below depend only on this sign structure, not on detailed tuning.

$$\Delta Q_{\text{gen}}(t) = -\frac{1}{K_Q} (V(t) - V_{\text{ref}}) S_{\text{base}}, \tag{4}$$

where V(t) is the local transmission-level voltage magnitude, V_{ref} the AVR/reference setpoint, K_Q the reactive-droop gain), and S_{base} the MVA base. Here $K_Q > 0$ and V is per unit, and the expression is a local proportional law and only used to establish the absorption/injection direction. The threshold model reflecting the P.O. 7.4 band can be written as

$$\Delta Q_{\rm gen}(t) \approx \begin{cases} 0, & 405 \leq V \leq 410 \; {\rm kV} \quad {\rm (no~obligation)} \\ + \Delta Q_{\rm lim}, & V < 405 \; {\rm kV} \quad {\rm (inject~to~raise} \; V) \\ - \Delta Q_{\rm lim}, & V > 410 \; {\rm kV} \quad {\rm (absorb~to~lower} \; V), \end{cases} \tag{5}$$

with $\Delta Q_{\rm lim} > 0$ a capability/limit term. Within the narrow no-obligation band around the setpoint, the net dynamic Q response is effectively negligible on the onset timescale, outside the band, proportional action with standard rate limits applies until capability is reached. This means that 405-410 kV is a no-requirement zone, above 410 kV units must absorb Q, and below 405 kV units must inject Q. In practice, P.O. 7.4 performance is evaluated over hourly windows using an adequacy criterion (e.g., requiring that at least a specified fraction of SCADA samples within the hour meet the voltage-control obligation). This percentile-style assessment can dilute very short excursions and hence does not enforce sub-second reactive response during the onset of a cascade. During the cascade, this meant that dynamic compensation was negligible within the voltage range where the cascade initiated (405–420 kV). Consequently, the algebraic constraint (3b) rebalanced entirely through voltage rise. Each subsequent trip removed additional reactive absorption, forcing progressively larger voltage increases in an unstable positive feedback loop.

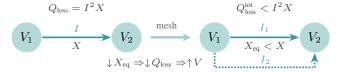


Figure 5: Meshing/energization mechanism—a simple illustration.

To quantify the voltage sensitivity, consider the linearized reactive power-voltage relationship from the power flow Jacobian embedded:

$$\Delta Q(t) = J_{QV} \Delta V(t) + J_{Q\theta} \Delta \theta(t), \tag{6}$$

where $J_{QV} = \partial Q/\partial V$ and $J_{Q\theta} = \partial Q/\partial \theta$ are the reactive-power Jacobian blocks evaluated at the operating point underlying (2). Here J_{QV} is the Q-V block associated with the algebraic matrix A_{aa} , while the coupling of angle changes appears through $J_{Q\theta}$. If angle deviations are modest over the first seconds (so $\Delta \theta(t) \approx 0$ because electromechanical angles evolve on slower (inertial) time scales, while voltages react quasi-instantaneously to ΔQ through the algebraic network), a net reactive deficit $\Delta Q_{\rm net} < 0$ implies a local voltage rise $\Delta V(t) \approx -J_{QV}^{-1} \Delta Q(t)$, unless dynamic compensation from (4) is promptly provided. Note that the inverse map is only invoked at the pre-event point, proximity to a voltage-stability boundary would render J_{OV} ill-conditioned and the linear map unreliable. For the onset window, $\Delta \theta$ is of higher order relative to algebraic voltage changes, so the leading-order relation $\Delta V \approx -J_{QV}^{-1} \Delta Q$ is appropriate. This is precisely why, within the 405-420 kV range where (5) yields little or no action, the algebraic constraints rebalance primarily via voltage increase, and each subsequent trip (further reducing absorption) amplifies the rise.

C. Topology-Induced Amplification of Voltage Sensitivity

The operator actions taken between 12:07 and 12:20 to dampen inter-area oscillations fundamentally restructured the network's voltage-reactive power coupling by modifying the algebraic Jacobian A_{aa} . This restructuing of the Jacobian usually improves the conditioning of the system; unfortunately, this was not the case in this incident. In short, the remedial operator actions of topological changes inadvertently caused more harm than good.[‡] In what follows, we delineate the three topology changes that merit analysis from a voltage control perspective.

Network meshing. Meshing several 400 kV circuits to damp morning oscillations created parallel paths that reduced equivalent impedance. For two parallel lines with reactances X_1 and X_2 , the equivalent reactance is $X_{\rm eq} = X_1 X_2/(X_1 + X_2) < \min(X_1, X_2)$. This modification propagates through the admittance matrix \boldsymbol{Y} in \boldsymbol{A}_{aa} . In the nodal admittance matrix \boldsymbol{Y} , adding a parallel branch increases the affected susceptance entries (and the connected diagonal terms) by $jB_{\rm eq}$ with $B_{\rm eq} = -1/X_{\rm eq}$. The meshing mechanism and its impact on reactive power and voltages is succinctly illustrated in Fig. 5.

The reduced impedance decreased reactive power consumption in transmission lines. Since $Q_{\text{line}} = |I|^2 X$, and current magnitude $|I| = |V_i - V_j|/|Z|$, parallel paths redistribute current such that aggregate reactive losses decrease (under constant power transfer):

aggregate reactive losses decrease (under constant power transfer):
$$Q_{\rm loss}^{\rm total}(t) = \sum_{(i,j)} \left| I_{ij}(t) \right|^2 X_{ij} < Q_{\rm loss}^{\rm original}(t), \tag{7}$$

[‡]All of these actions as documented in incident reports were taken in a matter of seconds and our analysis is using the power of hindsight which the operator did not have the luxury to afford at the time.

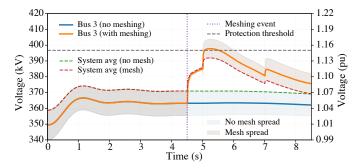


Figure 6: Impact of meshing on overvoltage: energizing parallel 345-kV circuits lowers the effective series reactance and natural MVAr absorption, steepening the Q-V response. The meshed case (red) climbs higher than the non-meshed case (teal); dashed lines are system averages and shading shows deviation from those averages.

This removed a natural voltage regulation mechanism: in the original topology, voltage increases drive higher line currents that consume more reactive power, providing negative feedback. The meshed topology weakened this self-regulating effect. Operationally, remedial meshing is performed while maintaining (or reducing) corridor MW transfers and regulating boundary voltages with existing controls. Under these conditions—same or lower active transfer and regulated terminal voltages—lower series reactance reduces natural series MVAr absorption and stiffens the line, so bus voltages rise for a given net Q deficit. By contrast, holding the angle difference' refers to a theoretical constraint in which δ is fixed as X changes, this is not the operational regime for oscillation mitigation and is noted only to clarify dependence on what is constrained.

Moreover, tighter electrical coupling can increase network stiffness in certain cases. At the operating point, it can increase the magnitude of cross-couplings in the local sensitivity from ΔQ to ΔV (i.e., entries of J_{QV}^{-1}), so a given disturbance at one location can induce larger voltage deviations at distant buses. Formally, a reactive power disturbance ΔQ_k at bus k produces voltage deviations:

$$\Delta V_i(t) \approx \sum_k \left[\boldsymbol{J}_{QV}^{-1}(t_0) \right]_{ik} \Delta Q_k(t), \tag{8}$$
 As seen in Fig. 6, energizing the parallel 345-kV circuit stiffens the

As seen in Fig. 6, energizing the parallel 345-kV circuit stiffens the line, after the meshing command at $t\!\approx\!4.5$ s, the Bus 3 trajectory (orange) rises above the 1.15 pu protection limit while the unmeshed case (blue) remains near the system average. The shaded envelopes quantify the voltage spread about the average for each operating mode, highlighting how reduced series reactance and diminished native MVAr absorption push the meshed corridor into overvoltage.

Shunt reactor disconnection. Throughout the morning, operators disconnected shunt reactors to counteract transient undervoltages during oscillations. Reactors provide controllable reactive absorption $Q_{\rm reactor} = -B_{\rm sh}V^2$, where $B_{\rm sh}$ is the susceptance. Removing reactors eliminated a fast-acting voltage control tool and reduced the system's aggregate reactive absorption capability by several hundred MVArs. More critically, it removed voltage-dependent absorption that provides natural damping:

$$\left. \frac{\partial Q_{\text{reactor}}(t)}{\partial V(t)} \right|_t = -2B_{\text{sh}}V(t) < 0.$$
 (9)

This negative derivative stabilizes voltage, when voltage rises, reactor absorption increases, pulling voltage back down. Disconnecting reactors eliminated this term from J_{QV} , further

degrading natural voltage regulation. This local derivative pertains to bus-shunt representation, placement behind series impedance reduces the local Jacobian entry but retains negative sign, so disconnecting the reactor still weakens voltage damping.

HVDC mode reconfiguration. The Spain-France HVDC link was switched from AC-emulation mode to fixed-power mode during oscillation mitigation. In AC-emulation mode, the HVDC link controls active power as a function of the voltage-angle difference between its terminals, providing inherent damping of inter-area oscillations. The reference active power is determined by [2, p. 88]:

$$T\dot{P}_{\text{set}}(t) + \left(P_{\text{set}}(t) - P_0\right) = K\left(\delta_A(t) - \delta_B(t)\right),\tag{10}$$

where P_0 is the base power setpoint (typically zero in emulation mode), K is the proportional gain, T is the phase-compensation time constant, and $\delta_A - \delta_B$ is the instantaneous voltage-angle difference between terminals (Spain and France). This control emulates the inertial and damping behavior of an AC interconnection, enabling the HVDC to modulate power flow in response to electromechanical oscillations. Fixed-power mode eliminates this response: $P_{\text{HVDC}}(t) = P_{\text{ref}}$ regardless of system conditions. When the cascade began, the HVDC maintained constant 1,000 MW export rather than adapting to the growing Iberian power deficit. In AC-emulation mode, the expanding angle difference would have provided damping and tended to reduce export given growing angle separation. Instead, the fixed setpoint forced Spain to continue exporting northward even as internal generation collapsed, accelerating frequency decline and earlier activation of underfrequency load shedding. The HVDC mode switch effectively removed a degree of freedom from the control vector u(t) in (2) at the moment when control flexibility was most needed, thereby making the system less controllable.

These three topology modifications collectively reshaped A_{aa} in ways that: (i) reduced natural reactive damping, (ii) amplified spatial propagation of voltage disturbances, and (iii) decreased available control authority, all while existing operational tools did not quantify the short-horizon voltage-control consequences of these actions.

D. Transmission-Distribution Observability Mismatch

A critical vulnerability emerged from the hierarchical voltage transformation at collector substations in distribution systems—or at the interface of transmission and distribution. The algebraic states $\boldsymbol{x}_a(t)$ in (1b) include voltages at all network buses. However, the system operator observes only transmission-level measurements $\boldsymbol{y}(t) = \boldsymbol{C}\boldsymbol{x}_a(t) + \boldsymbol{v}(t)$, where \boldsymbol{C} is the measurement matrix (indicating where PMUs are installed) and $\boldsymbol{v}(t)$ represents noise. Collector substations downstream of generation transformers constitute unobservable states: $[\boldsymbol{C}]_{ik} = 0$ for collector bus k. We note that ENTSO-E states that the grid lacks comprehensive high-resolution monitoring at these voltage levels [2, p. 115].

The voltage transformation relationship can be written as [7]:

$$V_{\text{coll}}(t) = \frac{V_{\text{trans}}(t)}{n_{\text{tap}}(t)} = \frac{V_{\text{trans}}(t)}{n_{\text{nom}}(1 + \alpha k_{\text{tap}}(t))}.$$
 (11)

where $n_{\rm nom}$ is the nominal turns ratio, α is the per-tap voltage adjustment, and $k_{\rm tap}$ is the tap position. With automatic tap changers, $k_{\rm tap}(t)$ adjusts to regulate $V_{\rm coll}$. Over the onset window, $k_{\rm tap}$ remains fixed by intentional delay, the transformer relation is therefore algebraic with a constant ratio, and secondary-side overvoltage may occur even when transmission-side measurements remain

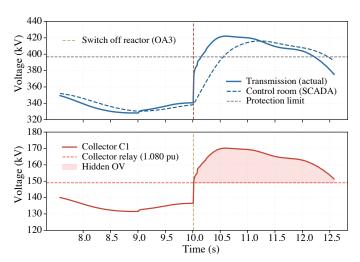


Figure 7: Protection-coordination mismatch: fixed GSU taps and collector voltage offsets bias the relay measurement

within limits. Tap changers operate with intentional delays to avoid *hunting*. During the morning's voltage transients, tap positions evolved to compensate for earlier undervoltages. By 12:32, transformers operated at tap positions optimized for lower transmission voltages, and could not respond quickly enough when voltages subsequently rose [3, p. 8].

When topology changes elevated transmission voltages to 418 kV at 12:32:57, the collector voltage became: $V_{\rm coll} = 418 \ {\rm kV}/n_{\rm tap} ({\rm set~for~}400 \ {\rm kV}) \approx 242 \ {\rm kV}.$ This exceeded the transformer overvoltage protection threshold (typically 1.10 pu), triggering the initial trip despite transmission voltage remaining within the 375–435 kV operational band.

From a power system DAE perspective, the collector voltage is an algebraic state coupled to the transmission system through the nonlinear transformer model in (1b). The critical insight here is that operator actions that modified A_{aa} to damp oscillations simultaneously affected unobservable algebraic states at collector level. The meshing actions increased transmission voltages through the mechanisms in (8), but operators could not observe the resulting violations of collector-level protection constraints. This created a feedback loop where each trip elevated voltages further, triggering additional collector-level trips that remained invisible until aggregate generation loss became severe. As seen in Fig. 7, the relay at collector C1 operates after a reactor switch off action at t = 10seven though the 345-kV transmission bus and the SCADA estimate remain below the 1.15 pu system limit. The upper panel shows the actual and SCADA transmission voltages and the protection limit, the lower panel shows the 138-kV collector voltage and its 1.08 pu relay setting. With a fixed GSU tap, the low-side voltage runs several percent higher in pu than the high-side reading, so the collector relay enters the shaded hidden overvoltage region and trips while the control room still sees an acceptable voltage.

The observability gap is further compounded by distributed generation operating in reverse power flow mode. Traditional grid design assumes distribution networks are reactive sinks: $Q_{\rm dist} < 0$ (absorp-

§Hunting is rapid, oscillatory tap switching when small voltage swings or noise make an On-Load Tap Changer (OLTC) overshoot its target. Because taps are discrete, the controller can bounce between steps; a deadband and time delay are added so a deviation must persist before a tap change is made.

tion). On April 28, according to REE, distribution networks injected approximately 760 MVAr into transmission (AA1), which altered the voltage control problem. The distribution network now acts as a voltage source with poorly regulated output, and transmission operators lack visibility into its state. The cascade mechanism relied on having generation at collector voltage levels: each transformer trip removed hundreds of MVAr of reactive absorption, driving further voltage rise. Had these been traditional load buses, voltage increases would have elicited stabilizing reactive consumption rather than destabilizing absorption loss through protection-triggered disconnections. This architectural mismatch, where a transmission system is designed for one-way hierarchical control operating under bidirectional power flows with thousands of uncoordinated sources, represents a vulnerability that voltage stability tools must address.

E. Coupled Voltage-Frequency Dynamics and the UFLS Paradox

The cascade exhibited bidirectional coupling between voltage and frequency instabilities, visible in the interaction between differential states $x_d(t)$ (such as frequencies) and algebraic states $x_a(t)$ (voltages) in (2). For synchronous generators, the swing equation governing frequency deviation $\Delta\omega_i(t) = \omega_i(t) - \omega_0$ is:

$$2H_i\dot{\omega}_i(t) = P_{m,i}(t) - P_{e,i}(\boldsymbol{v}(t),\boldsymbol{\theta}(t)) - D_i\Delta\omega_i(t), \quad (12)$$

where H_i is the inertia constant and electrical power $P_{e,i}$ depends on voltage magnitudes and angles:

$$P_{e,i}(t) = \sum_{j} v_i(t)v_j(t)[G_{ij}\cos\theta_{ij}(t) + B_{ij}\sin\theta_{ij}(t)]$$
 (13)

This creates the coupling: voltage magnitudes (algebraic states) affect power extraction $P_{e,i}(t)$ and hence frequency dynamics $\dot{\omega}_i(t)$, while angle deviations $\Delta\theta_{ij}(t)$ driven by frequency imbalance modify reactive power flows through the $\sin\theta_{ij}(t)$ terms in (3b).

When generation trips removed both active and reactive power, the system experienced simultaneous deficits. The active deficit created power imbalance in (12), driving frequency decline. The reactive deficit forced voltage rise through (3b). These instabilities reinforced each other: at the pre-event operating point, the relevant transfer corridors satisfy $\partial P_e/\partial V>0$, so higher voltages increase electrical power extraction and reduce frequency support margin locally, the sign may differ away from this point and does not affect the voltage-driven initiation analyzed here. All such sign statements are strictly local to the pre-event operating point used for linearization. Conversely, as frequency dropped and angles deviated, the $\sin\theta_{ij}$ terms in (3b) modified reactive flows in ways that could exacerbate voltage rise in certain network regions.

The most counterintuitive coupling emerged through underfrequency load shedding (UFLS). At 12:33:19, with frequency below 49 Hz, UFLS activated to arrest frequency decline by removing load (AA2). However, load shedding simultaneously removes reactive consumption, which according to (3b) forces voltage increase:

$$\Delta V(t) \propto - {\bf J}_{OV}^{-1}(t_0) \Delta Q_{\rm load}(t) > 0 \quad {\rm when} \ \Delta Q_{\rm load}(t) < 0, \ \ (14)$$

This created a paradox: the protection mechanism designed to stabilize frequency destabilized voltage. Load shedding removed large blocks of loads and several hundred MVAr of reactive consumption, which REE's report confirms caused further voltage increases throughout the system. The UFLS action was procedurally correct for frequency stability but accelerated the voltage cascade, a perverse interaction that is visible only through the coupled DAE structure (2).

At 12:33:21, loss of synchronism occurred at f = 48.46 Hz when angular separation exceeded stability limits: $\max_{i,j} |\delta_i - \delta_j| > \delta_{\rm crit}$. The AC interconnections correctly tripped to prevent instability propagation to Continental Europe. The paradox is that while frequency collapse was the proximate cause of islanding, the root instability was voltage-driven: had adequate reactive response been available to contain that initial voltage rise, the subsequent protection-triggered disconnections—and the ensuing frequency and synchronism collapse—would likely have been avoided.

VI.

REPLICATION RESULTS: CASCADE ONSET AND PROGRESSION

The replication framework described in Section III implements the physics-based cascade model using ANDES [5], an open-source DAE solver for power system dynamics. The complete simulation code, including IEEE 39-bus case setup, collector bus configuration, and protection logic, is available on GitHub [1]. While the model structure is conceptually straightforward (explicit collector buses with fixed taps, real device trips computed via standard power flow), practical implementation involves non-trivial calibration of tap positions, protection thresholds, and device ratings to match the voltage sensitivities observed in incident reports.

We replicate the cascade timeline by executing operator actions (meshing, reactor removal, export reduction, etc) at times matching the incident chronology, then monitoring collector voltages for protection threshold violations. When a collector exceeds its threshold for the configured dwell time, the associated generation and reactive devices disconnect, and ANDES computes the resulting voltage response. Fig. 8 plots the replicated cascade on the IEEE-39 system, showing the monitored 345-kV buses (3, 4,16, 19, 20, 21, 22, 33, 35, 38). Operator actions between t = 4.5 and 10s—closing four 345-kV parallels, opening both shunt reactors, switching the HVDC link to fixed export, and easing exports—gradually stiffen the corridor and remove roughly 200 MVAr of controllable absorption, so transmission voltages drift upward even though they initially remain below the 1.15 pu limit. When collector C1 trips at t = 10.08s (shedding 78 MVAr), two neighbouring collectors (C4 and C3) follow within 0.06s, shedding an additional 126 MVAr. Each trip forces algebraic rebalancing per (3b), raising voltages at adjacent GSUs and accelerating the next trip. The pink window ($t \approx 10.1$ -13.4s) highlights this rapid positive-feedback phase, matching the ENTSO-E report's description of several 400-kV substations disconnected within seconds. UFLS arms at t = 11.36s, and collapse is detected at t = 13.44s.

Fig. 9 quantifies the cascade driver by plotting maximum 345-kV transmission voltage (red, left axis) against cumulative reactive absorption lost (blue, right axis). During cascade onset (t=4.5–11 s), each operator action and collector trip removes reactive absorption, and the algebraic network constraints force voltage rise per (3b). The first three collector trips remove approximately 460 MVAr cumulatively and drive peak voltage from 370 kV to 433 kV, breaching the 387 kV protection threshold. The tight $Q{\to}V$ coupling is obvious, each increase in lost MVAr (blue) immediately elevates transmission voltage (red), with the relationship quantified by $\Delta V \approx -J_{QV}^{-1}\Delta Q_{\rm abs}$. After UFLS activation at $t\approx$ 11 s, voltage declines despite continued reactive losses—marking the transition where coupled frequency and angular dynamics dominate over the precipitating reactive imbalance. The figure thus distinguishes the cascade

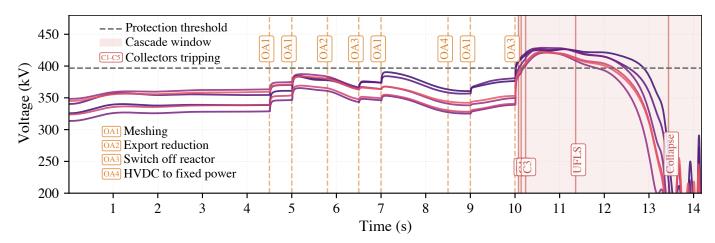


Figure 8: Replicated event sequence and voltage response on IEEE 39-bus system: selected 345-kV bus voltages versus time with the 1.12 pu protection limit (gray dashed). Labels mark operator actions and automatic actions. C1–C5 denote collector-level trips. The pink band shows the cascade window starting at the first collector trip.

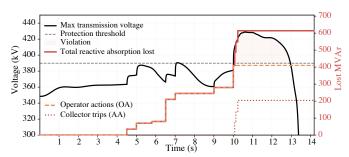


Figure 9: Reactive power-voltage coupling during cascade. Maximum transmission voltage (left axis) rises as cumulative reactive absorption is lost through collector trips (right axis). Q-V coupling drives voltage increase during cascade (t=4.9-13.3 s).

initiation mechanism (reactive-driven voltage rise) from the termination mechanism (coupled instability collapse). Additional diagnostic plots and case studies are provided in the GitHub repo [1, figures.py], including collector-level voltage traces and two system-level summaries: a Q–V phase portrait that relates cumulative loss of reactive absorption to the system-average 345-kV voltage, and a headroom/exceedance plot that tracks both the remaining kV margin and the fraction of EHV buses already above their overvoltage threshold.

It is important to note that the measures taken by operators to dampen oscillations—extra meshing of 400 kV circuits, switching off reactors, and placing the HVDC in fixed—power mode—were correct and followed established protocols. While some of these actions may have decreased the system's ability to absorb reactive power, further raising voltages and narrowing voltage control margins, it remains unclear whether a different sequence of actions would have avoided a blackout, particularly given the low share of synchronous machines and their inadequate voltage regulation, and the fact that several trip causes and oscillation triggers remain unknown and the reports characterize the event as a multifactor overvoltage cascade under active investigation.

What is clear, however, is that the blackout stemmed from multiple interacting factors: systematic generator non-compliance with dynamic voltage control obligations, regulatory delays in updating grid codes despite warnings from REE as early as 2020, limited observability at subtransmission and collector voltage levels, and the fundamental challenge of operating a grid engineered for one-way, hierarchical control under distributed generation with reverse flows.

Addressing voltage collapse risk requires addressing these multiple dimensions through improved reactive capability, DAE system level understanding, enforced device compliance, expanded monitoring infrastructure, and updated operational procedures. Within this broader context, several operational gaps emerged: operators lacked practical methods to assess whether their topology changes would maintain or degrade voltage control margins, had limited visibility into collector-level conditions where protection operated, and possessed no tools to evaluate coupled voltage-frequency stability during emergency actions. These gaps, rooted in the mismatch between conventional operational paradigms and high-IBR system dynamics, motivate the practical mitigation measures presented next.

VII. PRACTICAL MEASURES AND CONCLUDING REMARKS

The analysis presented throughout this paper highlights several operational gaps that, while not individually catastrophic, converged to enable a cascading voltage collapse. The blackout required multiple improbable conditions occurring simultaneously: high renewable penetration reducing available dynamic reactive reserves, systematic generator non-compliance with voltage control obligations, protocol-driven topology modifications that inadvertently degraded voltage control, and collector-level overvoltages invisible to transmission operators.

Addressing one of these factors differently *might* have prevented the cascade. However, the incident reveals a deeper challenge in modern power systems: conventional operator decision frameworks developed for synchronous-dominated grids may inadequately account for voltage-reactive dynamics on sub-second to few-second timescales. Traditional protocols for oscillation damping—energizing parallel circuits, switching reactors, adjusting exports—correctly addressed inter-area modes but lacked tools to assess simultaneous impacts on voltage control margins. Next, we suggest practical measures to mitigate and prevent similar incidents.

A. Practical Measures

Addressing the aforementioned gaps requires augmenting existing operational procedures with decision support tools that enable operators to evaluate voltage-reactive consequences of topology changes before implementation. We identify four practical measures that could close the operational gaps exposed by the Iberian blackout:

Dynamic voltage margin verification. Before executing topology modifications, operators should verify that available reactive devices retain sufficient control authority to regulate voltages on timescales relevant for preventing protection-triggered trips. This requires computing voltage sensitivities that account for finite device response times, then solving a simple feasibility problem: can available reactive adjustments maintain all bus voltages within acceptable bounds given device rating constraints? Infeasibility signals inadequate control authority and flags the proposed topology change for reconsideration or compensating actions such as energizing shunt reactors before meshing. The assessment requires several sparse matrix factorizations which enable real-time evaluation within existing energy management systems by augmenting contingency analysis modules with voltage control screening.

Stability-aware meshing. While operators routinely verify that meshing improves inter-area damping through modal analysis, conventional eigenvalue screening assumes small deviations around equilibrium and cannot capture the nonlinear protection-triggered cascades observed during the blackout. A more robust approach computes voltage recovery margins using trajectory sensitivity analysis [8], for each candidate meshing configuration, simulate the system's response to standardized reactive power disturbances (emulating generator trips or capacitor bank switching), then quantify the maximum disturbance magnitude that the system can absorb while maintaining voltages within protection deadbands over a finite horizon. The approach accounts for nonlinear voltage dynamics, discrete protection actions, and finite device response times. Implementation can leverage methods from Lyapunov theory by constructing an energy function that bounds voltage trajectories, compute the distance from current operating point to the boundary of the stability region, and flag meshing configurations that reduce this distance below acceptable thresholds. Offline precomputation builds lookup tables mapping operating conditions and line pairs to recovery margins which enables rapid online screening. This reveals when impedance reduction improves angular stability at the cost of voltage control, a tradeoff invisible to modal analysis alone.

Voltage-aware emergency control coordination. Underfrequency load shedding stabilizes frequency but removes reactive consumption, paradoxically worsening overvoltage conditions. Emergency control schemes should coordinate frequency and voltage objectives by preferentially shedding loads with low reactive-to-active ratios during overvoltage situations, preserving reactive sinks when voltage margins are critical. A coupled vulnerability index quantifies combined frequency-voltage stress, enabling load selection that addresses the dominant instability mode. Implementation requires upgrading UFLS logic to include voltage measurements and adaptive optimization rather than fixed shedding schedules [9]. While increasing scheme complexity, modern digital relays and communication infrastructure make such coordination technically feasible and operationally manageable.

These measures are simple, computationally tractable and allow for real-time or near-real-time deployment, and can be integrated with existing workflows. None individually guarantees prevention of all voltage collapse events but collectively provide defense-in-depth by closing specific gaps that enabled the Iberian cascade. Implementing them will be the authors' focus in the next year or two, and beyond the scope of this paper.

B. Concluding Remarks

The Iberian Peninsula blackout resulted from an improbable confluence of adverse conditions occurring within a critical 30-second window. While this paper systematically analyzes the cascade mechanisms and contributing factors, our analysis benefits from hindsight and time, luxuries unavailable to operators managing real-time oscillations with incomplete information under established protocols. The objective of this paper is *not* to project blame; the objective is to offer an educationally rigorous analysis.

To summarize, the blackout stemmed from four converging factors: (i) systematic generator non-compliance with dynamic voltage control obligations, (ii) protocol-driven topology modifications (meshing, reactor switching, HVDC reconfiguration) that dampened frequency oscillations but inadvertently degraded voltage control margins, (iii) lagged transformer tap positions creating collector-level overvoltages invisible to transmission operators, and (iv) architectural mismatch between grid design assumptions and high-IBR operating reality. Addressing any single factor differently could likely have prevented the cascade. Attributing the blackout to high renewable penetration alone misses the mechanism: the cascade stemmed from inadequate voltage control response and limited observability, not from renewables per se. Two statements are simultaneously true: high renewable penetration (82%) reduced available dynamic reactive reserves and shaped system conditions, and a similar overvoltage cascade could occur at lower penetrations if the same control and observability gaps persist. Blaming renewables for this blackout is akin to blaming highway traffic for an accident caused by malfunctioning traffic signals—the congestion may have been a factor, but the root failure lay in control infrastructure inadequacy.

This raises the central question: was this event a *Black Swan* (rare, unpredictable, obvious only in hindsight) or a *Gray Rhino* (visible, high-impact threat approaching yet neglected)? The evidence suggests it was predominantly *a Gray Rhino* The vulnerabilities were known and documented: REE issued warnings about voltage control deficiencies in high-renewable scenarios as early as 2020 [10, p. 5]; generator non-compliance was unfortunately systematic; collector-level observability gaps were recognized architectural limitations; and ENTSO-E's post-2016 recommendations for dynamic assessments remained unimplemented [11, p. 4]. The system routinely operated near voltage control limits during high-renewable, light-load conditions. The threat was visible, probable, and somewhat palpable. Renewables did not cause the blackout; a sequence of unfortunate events and lack of heeding advice did.

However, the specific sequence proved improbable: a forced oscillation of uncertain origin, morning topology changes converging to deplete reactive reserves, precise timing of the first trip when all preconditions aligned, and sub-minute cascade acceleration through invisible mechanisms. The 0.63 Hz oscillation origin remains unresolved [2, p. 59] at the time of writing in November 2025. While the vulnerabilities were gray rhinos, their simultaneous manifestation retained black swan characteristics.

Beyond philosophical discrepancies, the practical implication are as follows. Addressing Gray Rhino risks through the measures can substantially reduce a similar blackout probability even if Black Swan timing uncertainties remain irreducible. Defense-in-depth ensures improbable combinations cannot cascade unchecked. The incident demonstrates that in high-IBR systems, voltage-reactive

dynamics on sub-second timescales demand operational attention historically reserved for frequency-active dynamics. The key insight from our DAE-based analysis is that voltage (algebraic states), frequency (differential states), and angles (differential states) are tightly coupled. Operator actions modified the algebraic coupling structure to damp oscillations without tools to assess how these changes propagated through the coupled DAE loops to degrade voltage controllability on cascade-relevant timescales. Conventional frameworks require augmentation to capture these multi-timescale, cross-domain couplings.

ACKNOWLEDGMENT

This work was birthed out of a discussion with Mads Almassalkhi during the 2025 American Control Conference in Denver, in which he highlighted the critical role of network meshing in the Iberian blackout. Concurrently, the authors were examining the influence of topological variations on grid stability and controllability. We acknowledge Mads for his insightful conversations and constructive feedback.

REFERENCES

- [1] A. A. Albustami, "Iberian_blackout: replication of the 28 april 2025 iberian blackout," https://github.com/abdallahbustami/Iberian_blackout, 2025.
- [2] ENTSO-E Expert Panel, "Iberian blackout of 28 april 2025: incident report," ENTSO-E, Tech. Rep., 2025. [Online]. Available: https://eepublicdownloads. blob.core.windows.net/public-cdn-container/clean-documents/Publications/ 2025/entso-e_incident_report_ES-PT_April_2025_06.pdf
- [3] Red Eléctrica de España, "28 april 2025 incident in the spanish peninsular system: report and recommendations," Red Eléctrica de España, Tech. Rep., 2025. [Online]. Available: https://dlnlo4zeyfu21r.cloudfront.net/WEB_ Incident_%2028A_SpanishPeninsularElectricalSystem_18june25.pdf
- [4] Comité 28-A, "Non-confidential report on the 28 april 2025 electricity crisis," Gobierno de España, Tech. Rep., 2025. [Online]. Available: https://www.lamoncloa.gob.es/consejodeministros/resumenes/Documents/ 2025/Informe-no-confidencial-Comite-de-analisis-28A.pdf
- [5] H. Cui, F. Li, and K. Tomsovic, "Hybrid symbolic-numeric framework for power system modeling and analysis," *IEEE Transactions on Power Systems*, vol. 36, no. 2, pp. 1373–1384, 2020.
- [6] P. W. Sauer, M. A. Pai, and J. H. Chow, Power system dynamics and stability: with synchrophasor measurement and power system toolbox. John Wiley & Sons. 2017.
- [7] P. Kundur et al., "Power system stability," Power system stability and control, vol. 10, no. 1, pp. 7–1, 2007.
- [8] G. Hou and V. Vittal, "Trajectory sensitivity based preventive control of voltage instability considering load uncertainties," *IEEE Transactions on Power Systems*, vol. 27, no. 4, pp. 2280–2288, 2012.
- [9] A. K. Singh and M. Fozdar, "Event-driven frequency and voltage stability predictive assessment and unified load shedding," *IET Generation*, *Transmission & Distribution*, vol. 13, no. 19, pp. 4410–4420, 2019.
- [10] Red Eléctrica de España, "Explanatory report on non-frequency services in the spanish peninsular system," Red Eléctrica de España, Tech. Rep., 2021. [Online]. Available: https://www.cnmc.es/sites/default/files/editor_contenidos/Energia/Consulta%20Publica/5_DCOOR_DE_009_20_ AnexoIII_Memoria_REE_prop_CondicionesSnF_24062021.pdf
- [11] ENTSO-E System Protection and Dynamics (SPD), "Inter-area oscillations in the continental europe synchronous area," ENTSO-E, Tech. Rep., 2017. [Online]. Available: https://eepublicdownloads.entsoe.eu/clean-documents/ SOC%20documents/Regional_Groups_Continental_Europe/2017/CE_ inter-area_oscillations_Dec_1st_2016_PUBLIC_V7.pdf