Supporting information

Electrodynamic Stommel framework: methods, evidence boundaries and validation plan

Mature V3 — integrated supplementary information — 14 November 2026.

Rolf Bastiaanssen, Bax & Company, Carrer de Casp 118, 08013 Barcelona, Spain. Correspondence: r.bastiaanssen@baxcompany.com. ORCID: https://orcid.org/0009-0000-6119-429X. The author declares no competing interests.

These notes distinguish implemented methods, numerical execution and proposed validation. They accompany the integrated Perspective; they do not establish matched physical accuracy.

Supplementary Note 1 — Claim ledger and novelty

Established background: power systems exhibit multiple stability phenomena and characteristic times; network coupling matters; non-normal dynamics may amplify perturbations; current limits constrain converters. None of these principles is claimed as newly discovered.

Implemented contribution: an interactive scale-oriented interpretation layer, common network-derived inputs for a constructed field, explicit machine-pair harmonic observables, mechanism-level incident paths, independent diagnostic panels, user-supplied recording analysis and auditable draft dossiers.

Candidate novelty: integration of those elements into a reproducible evidence-aware workflow. Establish priority through a systematic search of power-system scale maps, stability atlases, vulnerability visualizations, operator tools, digital twins and regulatory sandboxes. The present reference list is a starting point, not an exhaustive state-of-the-art review.

Not established: unique first-in-literature status; probability calibration; accurate prediction of historical trips; general current-limit-induced DAE-index change; universal penetration threshold; market-settleable shadow prices; 3–30-fold gains; or connection-queue capacity released.

Search protocol: record queries, dates, databases, deduplication and inclusion criteria. Compare against IEEE's extended classification, impedance/Nyquist workflows, RMS/EMT environments, spatial vulnerability tools and visual analytics. Evaluate whether existing approaches already implement the asserted combination. Two independent reviewers should resolve borderline inclusions and archive the comparison matrix.

Supplementary Note 2 — Coordinates and field construction

Let $x = \log_{10}(\tau/\mathrm{s}) \in [-5, 7]$ and $y = \log_{10}(L/\mathrm{m}) \in [0, 6]$. L denotes the assigned coupling/coherence extent unless a different metric is explicitly attached. Characteristic tau is not elapsed event time.

For a Gaussian footprint G_j, widths and centres are data conventions:

$$G_j(x, y) = \exp\left[-\frac{(x - c_{x,j})^2}{2s_{x,j}^2} - \frac{(y - c_{y,j})^2}{2s_{y,j}^2}\right]. \tag{1}$$

The plateau is

$$P(x, y) = 0.12 + 0.12 \exp\left[-\frac{(y - (x + 5)/2)^2}{2.5}\right] + 0.03 \sin(1.3x + y). \tag{2}$$

The exact implemented field, with eight bands, is

$$\Psi(x, y) = \max\left\{0, \frac{1}{8} \log\left(e^{8P} + \sum_{j=1}^{8} e^{8a_j G_j}\right) + C(x, y)\right\}. \tag{3}$$

The local RoCoF amplitude depends on y. The oscillation centre moves to log10(1/f_mode). C is an illustrative nonnegative weak-grid modulation, not a solved resonance surface:

$$C(x, y) = 1.3bG(x, y; -1.5, 3.5, 1.1, 1.2)[\sin(5.5x + 3.1y)\cos(2.3y - 1.7x)]^2 \tag{4}$$

when $b > 0.01$, otherwise zero.

The smoothing parameter is 8. For nine terms, $\max_j t_j \le \frac{1}{8}\log\sum_j e^{8t_j} \le \max_j t_j + \frac{\log 9}{8}$. The additive bound is approximately 0.27465. This is why an aggregate contour cannot be interpreted as an exact union of engineering limit exceedances.

Band centres (log seconds, log metres): EMT (-4,1); PLL (-2,3); RoCoF (-0.5,4.5); nadir (0.6,5.5); oscillation (log10(1/f),5.6); voltage (1.2,4.6); governor (2,5.6); adequacy (6,5.7). Gaussian widths are conventions in BAND_GEOMETRY. They are not observed universal process domains.

Examples of amplitudes: RoCoF uses $f_0\Delta P/(2H_{\mathrm{eff}}S)$, normalized by a scenario criterion and a local-inertia factor. PLL combines a generic converter-derived critical SCR with effective grid strength. Oscillation uses $0.0375/\max(\zeta, 0.005)$: its unit crossing is not identical to a 5% damping criterion. Negative damping is floored in this field amplitude and must remain visible in the original diagnostic.

No bound restricts Psi to [0,1]. No implementation equates it to an energy quotient or derives it from a resolvent. Non-normal and electrical diagnostics are separate analyses. The map's critical-area fraction is a sampled fraction of log-coordinate area. It is not failure probability or affected demand. Recompute at progressively finer grids and report maxima, area fractions and contours separately.

Operator counterpart: a diagnostic, not a fabricated envelope

For a frozen operating point p, specify an energy-coordinate operator A_E(p), input B_L and observation C_L. A scale-resolved transfer diagnostic can be defined as

$$R(\tau, L; p) = \| C_L \left(\mathrm{i}\frac{2\pi}{\tau} I - A_E(p) \right)^{-1} B_L \|_2 . \tag{5}$$

Here tau is forcing period in seconds, $\omega = 2\pi/\tau$. The implemented analytical companion now uses explicit balanced pair ports, not B=C=I. The full-state resolvent and glow remain separate. No independently derived continuous family A(L) exists. Distinct pair experiments are arranged by electrical extent without interpolation into a universal spatial response.

Conditional upper-bound identity. For stable A_E, the resolvent equals $\int_0^\infty e^{-\mathrm{i}\omega t} e^{A_E t}\,\mathrm{d}t$. If a independently certified semigroup estimate $\| e^{A_E t} \|_2 \le M e^{-at}$ holds for all t>=0, a>0, then $R \le \| C_L \|_2 \| B_L \|_2 M/a$ by the triangle inequality. This is a bound on that specified operator transfer. Sampled 0–10 s growth does not certify its tail, M or a. It provides no inequality relating the existing Gaussian Psi to R.

No envelope theorem is asserted. The review's proposed bounded spectral expression cannot be proved below the present Psi: its coefficients, normalization, L-dependent operator and calibration domain are unspecified. For a simple $A = -aI$, $\| (\mathrm{i}\omega I - A)^{-1} \| = 1/\sqrt{a^2+\omega^2}$ can be arbitrarily large near zero frequency as a approaches zero, while an unrelated constructed field stays finite. Even a clipped diagnostic can exceed an unconstrained visual index at some coordinates. A mathematical link requires extra assumptions, not rhetoric.

A future envelope study should freeze an admissible operator family, metric, B_L/C_L, frequency domain and normalization; fit any mapping on training cases; and report held-out violations and confidence intervals. Finite-grid success is an empirical domain-limited calibration, not a continuum conservative-envelope proof. The two layers currently have a provenance connection through selected network inputs, not a mathematical duality.

Proposed alternatives: normalized log-sum-exp; maximum of explicit normalized engineering margins; vector-valued margin display; uncertainty-aware contour bands; no scalar aggregation. Compare decision performance rather than merely selecting the most dramatic surface.

Implemented pair observable and normalization

For two inertia-bearing machine indices a,b, the physical disturbance pattern is $p = e_a - e_b$: +P GW at a and -P GW at b. Thus P is the power-transfer amplitude, not total net loss or two independent deficits. The measured output is frequency_a-frequency_b, not an inertia-weighted mean or absolute system frequency. With the retained basis U of Supplementary Note 4, $v = U^\top M^{-1/2}(e_a - e_b)$, $B = [0; v]$, $C = [0, v^\top/(2\pi)]$. Then $H = C(\mathrm{i}\omega I - A_E)^{-1}B$ and $G = |H|$ is Hz/GW. Output $Q = (0.1\,\mathrm{GW})G/(0.1\,\mathrm{Hz})$ is dimensionless, with declared display assumptions rather than a grid-code limit. Q=1 is reference-amplitude equality, not instability or a protection trigger. No peak-by-grid normalization is performed.

Port masks are explicit node IDs. Pair diameter is their effective-reactance distance divided by the common 400 kV circuit reactance-per-kilometre convention; it is not geographic metres. Enumerate unordered machine pairs, sort by extent and machine indices, then choose up to six evenly spaced rank positions including both endpoints.

Display discrete rows at those extents; no Gaussian spreading, continuous length interpolation or monotone-length-gain assumption. Frequencies are sampled at logarithmically spaced forcing periods 0.1–20 s (64 UI samples, 32 verification samples).

The operator's retained basis omits common-speed coupling, so the transfer concerns retained relative dynamics only. Balanced power does not guarantee exact full-DAE decoupling under heterogeneous damping; port reconstruction inherits that Galerkin approximation. Passive-bus injection is not silently mapped to a nearby machine. Networks with multiple connected components, fewer than two inertia-bearing nodes, negative symmetrized stiffness below $-10^{-7}\lambda_{\max}$, an empty retained basis or spectral abscissa $\geq -10^{-8}\,\mathrm{s}^{-1}$ return unavailable. A complex LU solve requires finite outputs and relative right-hand-side residual $\leq 10^{-7}$. That residual is a backward-consistency check, not a certified forward conditioning bound near a pole.

For stable oscillator $A = \left(\begin{smallmatrix} 0 & 1 \\ -\omega_n^2 & -d \end{smallmatrix}\right)$, $B = [0, 1/M]^\top$, $C = [0, 1/(2\pi)]$, and

$$G = \frac{\omega}{2\pi M\sqrt{(\omega_n^2 - \omega^2)^2 + (d\omega)^2}}. \tag{6}$$

Unit tests compare this closed form, coordinate sign invariance, input scaling, slow/fast limits and invalid ports. For regular stiffness the relative-frequency DC limit is zero; at high frequency G tends to zero. Large forced resonance is compatible with contractive lossless free evolution. No universal bound G<=1 is asserted.

Four IEEE39 parameter configurations pass independent SciPy lossless-transfer reconstruction. Nominal-resistance response is withheld because symmetrized stiffness is indefinite; these rejected cases remain in the evidence ledger rather than being relabelled as harmless rotations.

Supplementary Note 3 — Reduced network and electrical metric

The network is a structure-preserving linearized swing model with bus angles eliminated by Kron reduction. Machine internal angles couple through internal stiffness and corridor stiffness. The modal problem uses $M^{-1/2}K_{\mathrm{red}}M^{-1/2}$, with generic damping conventions. Real-system corridor capabilities and some coordinates are estimates or published aggregates; benchmark branches use series reactance and stored operating angles. The current IEEE39 residual-inertia convention creates 29 inertia-bearing nodes (28 nontrivial modes) at the baseline, not a full 10-machine model. Generation/load are normalized to 6.1 GW without re-solving the stored AC equilibrium. No per-grid frequency correction is allowed.

Generalized grid strength uses a reduced network admittance with multi-infeed scaling. This does not supply a vendor-specific stability threshold. Q–V analysis is a reduced reactive-voltage screen, not a full dynamic voltage-collapse model. Large-signal current-limited GFM checks supply a surviving-support fraction only through the common surfaceInputs helper.

For a positive-weight connected susceptance Laplacian B, $d_{ij} = (e_i - e_j)^\top B^+(e_i - e_j)$. The implementation uses fault-level branch weights bFault. Distances across islands are infinite. Equivalent kilometres divide per-unit effective reactance by 0.001875 on a 1 GVA, 400 kV single-circuit convention. This is neither geography nor an empirically inferred coherence radius.

Two-dimensional multidimensional scaling approximates a distance matrix and must publish stress. An exact resistance-distance Euclidean embedding naturally concerns $\sqrt{d}$; a display based on d is a separate approximation. Participation-weighted mode extent and clipped converter extents are further interpretive choices. Do not conceal those choices behind the word "topological".

Full operator derivation from a general DAE: for $\dot{x} = f(x, z, u)$, $0 = g(x, z, u)$, elimination gives $A = f_x - f_z g_z^{-1} g_x$ only when g_z is invertible and the assumptions hold. This identity is contextual theory, not proof that the application's aggregate field follows from every such DAE. Current limiting does not alone establish singularity of g_z or an index transition.

Supplementary Note 4 — Non-normal operator and norm

Let T=$M^{-1/2}K_{\mathrm{red}}M^{-1/2}$. Diagonalize sym $(T) = U\Lambda U^\top$, retaining positive eigenvalues above a relative threshold. Energy coordinates use $z = [\Lambda^{1/2}U^\top M^{1/2}\delta; U^\top M^{1/2}\dot{\delta}]$. The retained state matrix is

$$A_E = \begin{pmatrix} 0 & \Lambda^{1/2} \\ -(U^\top T U)\Lambda^{-1/2} & -D_r \end{pmatrix}. \tag{7}$$

Reported propagator gain is $g(t) = \| e^{A_E t} \|_2$. Quadratic-metric gain equals $g(t)^2$ in that specified metric; for nonsymmetric stiffness this positive metric is not asserted to equal physical stored energy. The lossless positive-energy model is contractive. This algebraic check is not empirical validation. The method excludes near-zero and nonpositive directions; disclose the excluded dimension and justify its physical interpretation.

COI projection, reference invariance and discarded directions
For n machines in one connected synchronous island, $M = \mathrm{diag}\,(M_i) > 0$, $q = M^{1/2}\mathbf{1}/\sqrt{\mathbf{1}^\top M \mathbf{1}}$, and $\Pi_\perp = I - qq^\top$. Thus $\Pi_\perp^2 = \Pi_\perp$, $\Pi_\perp^\top = \Pi_\perp$ and $\Pi_\perp q = 0$. In physical angle coordinates the corresponding projector is $P_\delta = I - \mathbf{1}(\mathbf{1}^\top M)/(\mathbf{1}^\top M \mathbf{1})$. The centre-of-inertia angle is $\theta_{\mathrm{coi}} = (\mathbf{1}^\top M \delta)/(\mathbf{1}^\top M \mathbf{1})$. Choose Q with $Q^\top Q = I$ and $QQ^\top = \Pi_\perp$; relative coordinates are $Q^\top M^{1/2}\delta$ and their velocities.

Reference-offset proof. Under $\delta \mapsto \delta + c\mathbf{1}$, the relative coordinates do not change because $Q^\top M^{1/2}\mathbf{1} = 0$. For a different orthonormal quotient basis $Q' = QO$, $O^\top O = I$, a fixed-energy-coordinate operator changes by orthogonal similarity $A' = S^\top A S$, $S = \mathrm{diag}\,(O, O)$. Then $e^{A't} = S^\top e^{At} S$ and $(zI - A')^{-1} = S^\top (zI - A)^{-1} S$. Their 2-norms, pseudospectra and continuous-time Kreiss quantity $\sup_{\mathrm{Re}\, z>0} \mathrm{Re}\, z \,\| (zI - A)^{-1} \|$ are identical. This proves invariance under basis/reference changes in a fixed metric, not under arbitrary nonorthogonal scaling or changes of metric.

Decoupling is conditional. If $Tq = 0$, write mass-weighted displacement as $Q\eta + q\theta$ and damping as $\bar{D} = M^{-1/2} D M^{-1/2}$. The relative acceleration contains $-Q^\top \bar{D} q \dot{\theta}$. Heterogeneous damping generally couples COI speed into relative dynamics. The lossy torque term can also feed relative motion into the COI equation through $q^\top T Q$. Removing COI angle does not justify removing COI speed. A time-varying reference rotation requires consistent treatment of its speed/acceleration. The application omits coupling to discarded coordinates; it does not simulate a governor enforcing a COI constraint. This Galerkin approximation omits that coupling; it is not an exact full-system quotient theorem.

For k connected synchronous islands, there are k rigid-angle directions when the corresponding Laplacian assumptions hold. The ideal lossless positive stiffness case has n-k positive relative-angle directions. The actual implementation instead retains eigenvectors of sym(T) above 1e-7*lambda_max, with dimension r; it removes n-r angle AND n-r speed directions. With a nonsymmetric lossy T, sym(T)'s truncated subspace need not equal the exact COI quotient. Near-zero or negative stiffness directions cannot be relabelled automatically as harmless rotations. Archive r, eigenvalues, threshold, island count and coupling residuals before claiming equivalence; no such diagnostic audit or Kreiss-constant calculation is claimed here.

Pseudospectrum is $\{z : \sigma_{\min}(zI - A) < \epsilon\}$. For a Hurwitz operator, a sampled imaginary-axis minimum is an approximate unstructured complex stability radius, not a guaranteed structured physical robustness margin. The implementation searches a finite frequency range and samples growth over 0–10 s at 0.05 s intervals. Approximate spectral norms and finite grids require convergence checks.

The glow compares $\rho_{\mathrm{lossy}}/\rho_{\mathrm{lossless}}$, where $\rho = \mathrm{dist}\,(\mathrm{i}\omega, \mathrm{spec}\,(A))/\sigma_{\min}(\mathrm{i}\omega I - A)$. The display clips the lossy/lossless excess ratio below at one, although that ratio need not exceed one mathematically. Its mapping to $\tau = 2\pi/\omega$ and L follows dominant-mode extents. Finite inverse/power iterations and frequency/time grids supply no certified numerical error bound. A bright patch is an explanatory display, not a calibrated warning probability. Different resistance settings may use different energy metrics; fixed-metric comparisons are needed for ST-5. Capability-class R/X values are shared estimates, not measured line resistances.

Supplementary Note 5 — Hybrid converter screen and rupture

The generic GFM model uses a VSM/droop-style local dynamic model and a current limiter, default Imax=1.2 pu. Retained fault voltage, duration, grid strength, priority convention and horizon must be recorded. Vendor controls, PLL freeze and current-loop gains may be stored in a profile without being simulated. User-supplied profiles are experiments, not authenticated vendor replicas.

Saturation and loss of synchronism are different events. The display decomposes $\Psi_{\mathrm{actual}} = \Psi_{\mathrm{linear}} + \Delta\Psi_{\mathrm{line}} + \Delta\Psi_{\mathrm{gfm}}$. The line contribution changes topology while holding GFM survival at unity; the converter contribution changes survival on the changed topology. This telescoping difference is exact for the sampled fields even when the field aggregation is nonlinear.

Explicit hybrid formulation and regularity test

Let h be a discrete mode in {voltage-source, current-limited, tripped}; x denotes controller/dynamic states and v algebraic network variables, not the plotted spatial coordinate. Within a mode: $\dot{x} = f_h(x, v, u), 0 = g_h(x, v, u)$. A voltage-source-to-limited guard is $\parallel i_{\mathrm{req}}(x, v, u) \parallel - I_{\max} = 0$ with an outward crossing; the reverse guard uses an inward crossing or a stated hysteresis margin. Protection transitions require a separately specified voltage/time, pole-slip or relay condition. A default radial limited output is $i_{\mathrm{out}} = I_{\max} i_{\mathrm{req}} / \parallel i_{\mathrm{req}} \parallel$ outside the current circle; d/q priorities are different active sets and must be recorded.

If a limiter changes output only, the continuous-state reset is $x^+ = x^-$; a trip applies its explicitly specified connection/state reset. Algebraic variables satisfy $g_{h^+}(x^+, v^+, u) = 0$ after switching, rather than being assumed continuous. A post-switch root must be obtained and checked. If $J_{v,h} = \partial g_h / \partial v$ is nonsingular at that root, the implicit-function theorem locally gives $v = \phi_h(x, u)$ and an index-1 semi-explicit representation for that mode. This is a local statement, not a global existence or uniqueness theorem. If J is singular, the test is inconclusive about index, solvability and impasse; further constraint analysis is required. Fixed reference/gauge constraints are necessary before testing a network Jacobian. Report $\sigma_{\min}(J)$, scaling and conditioning instead of relying on an unstable determinant test.

Counterexample to automatic index change. For the local algebraic equation $g(v) = v - v_{\mathrm{grid}} + Z i_{\mathrm{out}} = 0$ with prescribed limited current, $J_v = I$. Saturation can therefore coexist with a regular index-1 mode. If output depends on v, $J_v = I + Z \partial i_{\mathrm{out}} / \partial v$ and regularity must be checked at the operating solution. Neither example supplies the Jacobian of a complete vendor network model.

The implemented local screen integrates explicit reduced equations with algebraically evaluated currents; it does not construct a network DAE Jacobian, solve switching consistency or certify index changes. Its gfmRideThrough is a surviving support fraction under a deterministic screen, not a calibrated survival probability. Network regularity and saltation/switching stability remain proposed reference-model experiments. Radial clipping is continuous at the limit, but its Jacobian changes; "jump" need not mean a jump in terminal current or controller state.

Seams are masked where a contribution exceeds a declared plotting threshold. Their position is not a physical fracture coordinate or an inferred relay location. Within a fixed operating structure Psi remains smooth in x,y. A switching vector field may change continuously or discontinuously depending on the limiter; neither situation by itself proves a DAE-index change. Avoid "C0 fracture theorem" or "rank-drop proof" language.

Supplementary Note 6 — Historical inputs and reconstruction

Official report records are AEMO March 2017 (SA September 2016), National Grid ESO September 2019 (GB August 2019), and NERC/Texas RE September 2021 (Odessa May and June 2021). The current embedded numbers were previously transcribed without a complete page-locator audit. Report-labelled is not synonymous with independently verified.

Input ledger at package generation:

- SA: demand 1.826 GW report-tagged; non-synchronous proxy 0.77 derived and includes imports; reserve zero estimated. This share is not literal inverter-generation penetration.
- GB: demand 29 GW report-tagged; share 0.50 estimated; reserve 1 GW and LFDD volume 0.931 GW report-tagged.
- Odessa: demand 45 GW, share 0.35 and reserve 2.3 GW estimated. Prescribed trip comparator 1,112 MW is report-tagged, not predicted by the reduced model.

For every value, record source URL, exact report edition, page/figure, units, transcription author, second reviewer and ambiguity. Cross-check fault type and stage chronology. The supplied brief's Odessa SLG narrative does not match the current DLG replay convention. Do not change or selectively reinterpret either solely to make agreement appear stronger.

Prescribed staged deficits and outages drive a shared aggregate frequency equation:

$$\dot{\Delta f} = \frac{f_0}{2E_{\mathrm{island}}}[-P_{\mathrm{loss}}(t) + r + s - D\Delta f]. \tag{8}$$

$D = P_{\mathrm{load}}/f_0$; $3\dot{r} = \min(R, P_{\mathrm{sync}} \max(0, -\Delta f)/(0.04 f_0)) - r$. UFLS ramps a prescribed volume over 0.2 s. Integration uses 2 ms steps and a 50 fps synthetic trace. A frequency threshold at 0.94*f0 triggers a model collapse flag; later decline is scripted. This is not a post-collapse physical solution.

Illustrative regeneration: SA initial RoCoF approximately 5.03 Hz/s versus embedded 6.25 comparator; GB nadir 48.819 Hz and LFDD not triggered; Odessa model nadir approximately 59.755 Hz. Numerical coincidence is not proof of trip prediction. LFDD volume is input, not independent validation; the Odessa flag assesses a simplified sequence screen, not trip MW or device chronology.

Supplementary Note 7 — Measurements, COMTRADE and DMD

Decoder supports ASCII, signed-int16 BINARY, signed-int32 BINARY32 and IEEE FLOAT32; digital status words, calibration and timing remain explicit. Limits reject oversized or malformed inputs. Buffer-based analysis runs in a worker so the display remains interactive.

Sequence extraction requires synchronous A/B/C voltage channels and adequate samples per cycle. Missing or skewed channels must be disclosed. Frequency is used only when a measured channel exists; no frequency, inertia or spatial extent is fabricated from a filename. Raw waveforms can contain mains-frequency carrier and harmonics.

The DMD path uses truncated SVD and bounded overlapping-window recomputation, not a proven incremental streaming estimator. Its snapshot and channel bounds limit memory and numerical cost. One-channel analysis uses delay coordinates and cannot identify a multi-bus spatial shape. Report rank, singular-value cut-off, timebase, window, fit residual and sensitivity. Mode frequency/damping is a best fit to the measurement window, not independently validated physical state eigenvalues.

TLS-DMD and forward-backward comparison

For equally timed paired, mean-centred snapshots X and Y, ordinary DMD assumes X exact and fits Y approximately A*X. Noise in both blocks is an errors-in-variables problem. TLS-DMD stacks $Z = [X; Y]$, computes its truncated SVD $Z = U_z S_z V_z^\top$, retains right subspace $V_{z,r}$ and projects both blocks using $P = V_{z,r} V_{z,r}^\top$. Set $X_{\mathrm{tls}} = XP$, $Y_{\mathrm{tls}} = YP$, then obtain the reduced DMD operator from the truncated SVD $X_{\mathrm{tls}} = USV^\top$ as $\tilde{A} = U^\top Y_{\mathrm{tls}} V S^{-1}$. The implementation uses the equivalent compact blocks $XV_{z,r}$ and $YV_{z,r}$ to avoid forming a large P. It reports the TLS rank and the prediction residual against the original unprojected snapshots, not only the denoised fit.

TLS is now available in the shared DMD module and used by recording windows; ordinary DMD remains available for sensitivity comparisons. This corrects the asymmetric fit under an appropriate approximately isotropic errors-in-variables model. It does not eliminate coloured noise, delay-coordinate correlations, sensor skew, nonstationarity or model mismatch. Channels with different units or noise variance need a justified scaling/whitening protocol. No claim of unbiased field damping, guaranteed no ringing or physical-mode identification is made.

For an optional forward-backward estimator in the SAME reduced coordinate basis, fit $A_f = YX^+$ and $A_b = XY^+$, and form $A_{\mathrm{fb}} = (A_f A_b^{-1})^{1/2}$. In the noiseless invertible case both fits give the same A and its inverse, so the expression reduces to a branch of $\sqrt{A^2}$. The square-root branch must be chosen continuously and consistently with the sampled dynamics; singular backward operators, negative/complex branches and aliasing require explicit handling. FB-DMD is described as a proposed cross-check, not implemented or a universal cure for correlated noise.

Validation plan: use known damped oscillators plus independently generated white and coloured sensor noise over repeated seeds, report frequency/growth errors versus noise-free truth, compare OLS/TLS and optionally FB, sweep rank/window/channel scaling, and reserve recordings for held-out tests. A clean synthetic recovery test is software verification only. Methods references: Hemati et al., Theoretical and Computational Fluid Dynamics 31, 349–368 (2017), doi:10.1007/s00162-017-0432-2; Dawson et al., Experiments in Fluids 57, 42 (2016), doi:10.1007/s00348-016-2127-7. Confirm bibliographic records before submission.

Principal-log frequencies are aliased modulo sample rate; a negative real discrete eigenvalue lies at the Nyquist branch boundary, not an ordinary resolved complex spatial mode. A future recording study should use hold-out windows, independent channels, known events and reference mode estimates. Archive channel mapping and permission constraints. Compare inferred frequencies with established methods and quantify variation over windows rather than selecting an attractive mode overlay.

Supplementary Note 8 — Validation plan

The protocol in Supplementary Note 13 defines C0–C3 and ST-1–ST-5. Proposed gates are frequency relative error 5%; damping absolute error 0.02; VUF relative error 6%; CCT relative error 8%; and energy-gain relative error 10%. These are design choices supplied in the brief, not attained performance.

The executable comparator rejects mismatched metadata, missing solver identity, non-finite metrics, inappropriate zero denominators and unexecuted reference data. All twenty case/test pairs begin pending. Its numerical gate tests validate the comparator, not the power-system model. An unmatched IEEE39 baseline remains separately labelled.

Open-reference execution is now supplied in run_open_solver.py. NUM-C0–NUM-C3 independently assemble corridor stiffness and Kron elimination from prepared parameters in SciPy, never consuming TypeScript reduced matrices or eigenpairs as physical inputs. Frequency and shape matching uses MAC and assignment; the same reported PSS damping convention is reproduced, not independently validated. Nominal R/X harmonic responses are withheld; lossless pair gains are independently reconstructed from physical masks and transformations. Results and numerical tolerances are in Supplementary Note 12.

ANDES 2.0.0 executes its bundled ieee39_full.xlsx with ten GENROU machines. A positive-sequence bus-14 fault begins at 1 s, rf=0, xf=0.001 pu, clears after 100 ms and trips line 14–15. Runs through 8 s at nominal 5 and 2.5 ms trapezoidal steps archive rotor-frequency and bus-voltage traces, state matrices, spectra, deck bytes and SHA256, versions and warnings. Step shrinkage is permitted on convergence difficulty; actual timestamps are archived. The stock deck reports six positive-real eigenvalues and typical-limit/parameter-correction warnings, plus an ill-conditioned eigenvector calculation. Successful finite-horizon integration does not establish a stable reference equilibrium or recovery. This is an executed exploratory RMS reference, not a matched physical validation. Neither this run nor numerical verification satisfies ST-1–ST-5 or unbalanced EMT/vendor-current-limit validation.

Supplementary Note 9 — Locational stability screening and needs assessment

The policy modules compose existing physics, but economics are illustrative. Dossier fields preserve assumptions and indicate draft status. Screening thresholds are indicative or user-defined, not certification.

The welfare curve samples network-output perturbations rather than solving each stochastic network realization. Its simulated shedding fraction is model-conditional. EENS-like arithmetic multiplies simulated mean shedding by restoration hours and assumed faults per year; monetization multiplies by assumed VoLL. Curtailment uses a stylized annual factor and replacement-energy value, not chronological dispatch. Omitting capital costs and system feedback precludes a social-welfare claim.

Nodal DeltaPsi is a finite difference in global maximum field value per added 100 MW GFM, with global conventions unchanged. It is not dPsi/dH, a dollar price or a settlement-ready shadow value. Nonsmooth changes in the identity of the maximum can affect sensitivity. Test multiple capacity increments and track the maximizing location.

A policy pilot should specify authority, bounded capacity, study requirements, protection conditions, monitoring, stop criteria, sunset, liabilities and independent evaluation. The draft dossier does not grant a derogation or certify a grid-code result. Review legal wording with the relevant jurisdiction before external use.

Supplementary Note 10 — Narrative and user-study design

Conditional operating-region explorer

For a declared operating point p, topology T, disturbance set $\mathcal{D}$ and controls c, a genuinely secure operating region would require all relevant constraints $m_j(p, T, d, c) \geq 0$ for every $d \in \mathcal{D}$. The implemented explorer is a restricted sampled screen, not that complete region. It retains the current corridor states and contingency multiplier, samples eleven IBR proxy fractions including zero and one, and nine condenser capacities from zero to 1.6 times scenario peak demand in GVA. It plots finite positive generalised SCR on a logarithmic horizontal axis and IBR proxy on a linear vertical axis. Points without a finite strength coordinate are omitted from the plot, not from exported samples; no interpolated frontier or monotonicity theorem is asserted.

The GFM control is $\alpha = P_{\mathrm{GFM}}/(sS)$, with $P_{\mathrm{GFM}} = \alpha s S$ and s the IBR proxy. This is a generic capacity/support convention, not a matched fleet dispatch. Delivered condenser capacity and GFM quantity replace the corresponding Policy Lab quantities only inside this experiment; other interventions remain fixed. The shared full-delivery convention is selected explicitly. The network's generic GFM support treatment is not a vendor-specific representation of conversion from GFL to GFM.

Physical kinetic energy in the aggregate frequency convention is

$$E_{\mathrm{phys}} = [H_{\mathrm{full}}(1 - s) + H_{\mathrm{res}}]S + 2C, \tag{9}$$

with S in GW, C in GVA and the shared condenser inertia convention of 2 s. Equivalent energy is $E_{\mathrm{eq}} = E_{\mathrm{phys}} + H_{\mathrm{GFM}} P_{\mathrm{GFM}}$ in GW·s. Virtual support is not physically stored rotor energy or guaranteed immediate response. The initial frequency proxy is $r_0 = f_0 \mid \Delta P \mid /(2E_{\mathrm{eq}})$. This ignores latency, headroom, current limits, and local/noncoherent RoCoF. It is not the 500-ms nodal observable.

The four sampled conditions are: aggregate initial RoCoF below the scenario criterion; the shared nadir-band proxy at most one; least-damped inter-area $\zeta \geq 0.05$; and focus-network generalised SCR above the critical value obtained from the existing generic GFL model. The last is a screening comparison, not an exact multi-infeed eigenvalue-crossing theorem. It is inapplicable when no GFL remains. Absence of GFL cannot establish all-GFM stability. The combined result is the conjunction of these conditions. Frequency-only and GFL-strength-only displays expose why satisfying one mechanism does not imply satisfying another.

Q–V N–1 and current-limited fault simulations are omitted in this bounded interactive sweep and must be performed before operational use. A larger loss affects frequency margins without establishing PLL loss of synchronism. Condensers affect both physical inertia and strength in the shared model; the two quantities are not interchangeable. No fixed SCR range such as 1.5–2 establishes a universal boundary, and no new eponym or "phase-lock inversion frontier" is presented as an established scientific diagram. CSV exports include scenario, current topology, contingency, retained interventions, generic threshold, samples and the selected operating point. Verification checks units, GFM denominator, exclusion of virtual energy from physical inertia, deficit scaling, domain rejection and conjunction of criteria; physical validation remains pending.

The homepage is a sober incident-led investigation; the console is an academic dashboard. Both share model outputs. On narrow screens, an orthographic map avoids three-dimensional label collision while preserving the same field. Map rotations and familiar landmarks can aid orientation but risk concealing uncertainty behind apparent realism.

Evaluate three arms: conventional report/classification, scale map without Psi, and full map. Randomize order or use between-subject allocation to avoid learning carry-over. Recruit power engineers, regulators and executives separately. Pre-register comprehension, correct mechanism-scale distinction, suitable next-model selection, uncertainty recognition and false-confidence measures. Power analysis follows a pilot effect estimate; no arbitrary sample size or expected benefit is claimed.

Qualitative interviews should test interpretations of "stable", "probability", "electrical distance" and "tear". Successful communication means fewer consequential misunderstandings, not only faster task completion or higher enjoyment. An explanatory visualization can be harmful if it encourages causal certainty unsupported by the data.

Supplementary Note 11 — Reproducibility, governance and article conversion

Archive code commit, numerical snapshot, seeds, grid resolution, model assumptions and software versions. Preserve raw results and failures. Public deposit requires copyright, data and dependency licensing review. Author metadata and the supplied competing-interest declaration are recorded; no DOI, funding status or institutional endorsement is invented.

Before a research Article: establish prior art; formalize coordinate and margin definitions; complete an independent benchmark matrix; verify report inputs; obtain authorized recordings; quantify aggregation error; and test practical interpretation. A Nature Energy Perspective should remain a balanced conceptual argument, with original validation reserved for an appropriate research format. A Nature Communications Article needs substantive reproducible findings, not only a publication-quality interface.

Professional-journal material should focus on the bounded workflow and pilot needs, not recycle the same claimed original result. Disclose related submissions, avoid duplicate publication and secure editor permission for reused figures. The strongest outcome is a defensible chain from physical mechanism to next study and monitored decision.

Supplementary Note 12 — Frozen open-solver execution

Evidence generated 2026-11-06T14:52:56Z. Software: Python 3.13.15; NumPy 2.4.3; SciPy 1.18.0; ANDES 2.0.0. This edition reuses the frozen V1 evidence; it is not a new numerical run.

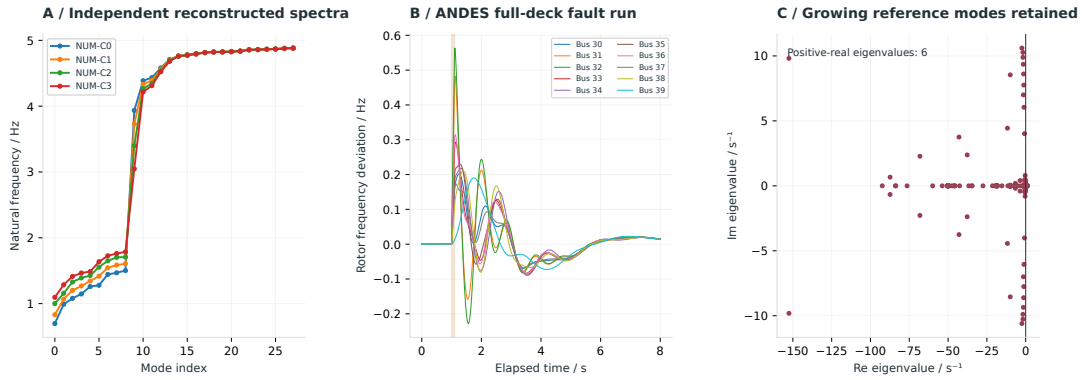
Reduced case	Modes	Frequency relative discrepancy	Minimum MAC	Pair-transfer scaled discrepancy
NUM-C0	28	5.73e-15	1.000000000000	3.33e-13
NUM-C1	28	8.72e-15	1.000000000000	6.34e-13
NUM-C2	28	5.09e-15	1.000000000000	2.44e-13
NUM-C3	28	5.68e-15	1.000000000000	6.01e-14

The maximum eigenpair residual is 2.16e-15. These are independent reconstructions of the same reduced equations, not physical validation or matched C0–C3 replacement cases. Damping-formula agreement checks the formula, not measured damping.

The unmatched ANDES deck has ten GENROU machines and 160 retained states. A bus-14 balanced fault at 1.0 s clears and trips line 14–15 at 1.1 s. Runs to 8 s use nominal 5 and 2.5 ms steps, with shrinking permitted. Maximum fine/coarse rotor-frequency difference is 0.000130738 Hz: time-step sensitivity, not error against truth.

ANDES reports 6 positive-real eigenvalues, spectral abscissa 1.03278 per second, parameter corrections and ill-conditioned eigenvectors. The input deck, raw traces, matrices, eigenvalues and full logs remain in the source/evidence archive. Stable equilibrium and recovery are not established. Nominal-resistance reduced IEEE39 response is withheld for indefinite symmetrized stiffness. **0/20 matched physical gates are complete.**

5 | Open execution does not erase model-equivalence gaps



Source: SciPy 1.18.0 and ANDES 2.0.0; 4 reduced numerical checks, NOT matched ST gates. Nominal-step rotor difference: 0.000131 Hz.
Reference controls/dispatch differ. Six growing modes, parameter corrections and ill-conditioning warnings unresolved; no stable-recovery or EMT validation.

Supplementary Figure S1 | Numerical reconstruction and an unmatched RMS execution. Spectra reconstructed from frozen parameters, ANDES rotor-frequency traces and reference eigenvalues. Growing modes are retained, not suppressed. Source: archived SciPy/ANDES execution; no EMT parity or stable-recovery validation.

Input SHA-256: `126be97269813411f84d44c01802a0013a9d654fe2830ca2437b34ac16d48356`. ANDES deck SHA-256: `9c2048dc94201ee48ffe816de65d53831fa1f4b0e7fdf32f24c50f1017edcfe5`.

Supplementary Note 13 — IEEE39 validation protocol

Purpose and status

This protocol tests a reduced screening framework against independently executed reference dynamics. None of its five gates has been attained. An exported IEEE39 baseline is not a fault simulation, nor independent validation. Freeze the protocol, observable definitions and input hashes before producing reference results. Tolerances below are proposed study design choices from the supplied brief, not journal requirements or industry standards.

Cases and operating points

Use a fully specified 39-bus, 10-generator data release with a branch inventory and solved AC operating point. Archive dispatch, load composition, machine dynamics, transformer grounding, converter controls, bases and initial conditions. The application has a benchmark-labelled network, but this does not establish equivalence to any particular PSCAD deck.

- C0: synchronous baseline, target 0% inverter injection.
- C1: target 40% dispatched inverter share; candidate replacements at buses 32–34.

- C2: target 70%; additional clustered replacements at buses 35–37.
- C3: target 85% inverter share, with GFM share explicitly reported as a fraction of total dispatch and of inverter capacity; candidate 100 MW GFM plant at bus 38.

Bus selections alone do not deliver those shares. Solve and document an allocation preserving feasibility and stated total generation. Do not silently interpret "85% IBR + 15% GFM" as 100% independent categories: GFM is a subset of inverter-based resources. Verify transformer-inclusive branch counts rather than asserting a universal corridor count. Initial voltage, power balance and controller equilibrium must agree before a dynamic comparison is valid.

Five stress tests

Test	Registered disturbance	Observable	Proposed gate
ST-1	3-phase bus-14 fault, 100 ms, clear line 14–15 if that branch exists in frozen deck	Matched inter-area frequency, identified over an agreed post-fault window	Relative error ≤5%
ST-2	Inertia sweep from 5.0 to 1.8 s with unchanged operating point and stated machine allocation	Matched mode damping ratio	Absolute difference ≤0.02
ST-3	SLG bus-36 fault; specify impedance, grounding and duration	VUF = abs(V2)/abs(V1), identical one-cycle estimation and 0–500 ms window	Relative error ≤6%
ST-4	Severe balanced dip; specify location, residual voltage and restoration	Critical clearing time with Imax=1.2 pu	Relative error ≤8%
ST-5	Multiply resistance of branch 16–17 by 3; re-solve operating point	Peak perturbation energy gain in a common registered metric	Relative error ≤10%

The existing app supports global R/X scaling, not a ready-made selective ST-5 branch run. Current sequence analysis is static: it cannot satisfy a 500 ms waveform parity claim without extending the experiment. The local GFM clearing-time routine is not a network-wide CCT prediction. Record unsupported comparisons as PENDING, not zero error.

Observable definitions

ST-1/2: remove COI mode, match frequency AND participation/MAC, use identical filtering and estimation windows; disclose modes that swap order. Estimate uncertainty using repeated window selection, not only one best-fit curve.

ST-3: export phase voltages at 10 kHz. Obtain positive/negative sequences with a common one-cycle window. Pre-register whether the scalar comparator receives peak VUF or windowed RMS VUF. Also publish the entire residual trace. Near-zero reference VUF requires a separately registered absolute criterion; never divide by zero or add an undisclosed epsilon. Do not claim the scalar gate validates PLL loss-of-lock.

ST-4: bracket clearing times, establish stable/unstable outcomes using the same horizon and angle/frequency/voltage criteria, then bisect. Report censored intervals (e.g. stable through maximum clearing time) rather than exact CCTs. The generic local model currently calls pole slip using its own criterion; harmonization is prerequisite, not a result.

ST-5: define E(t)=0.5$x(t)^TW^*x(t)$, W positive definite and fixed for both solvers after a documented state projection. Perturbations and initial energy must match. Compare Emax/E0. The application's gMax is an amplitude gain and must be squared for energy, with its metric and sampled time horizon disclosed. A single trajectory cannot establish an operator-norm maximum; use a registered perturbation ensemble or optimized initial condition. Exclude unobservable or unstable discarded directions only with justification.

Boundary-of-validity study: ST-3 and ST-4

ST-3 mismatch mechanisms: a positive-sequence Kron reduction omits explicit negative/zero-sequence dynamics, grounding/transformer detail and waveform/control transients. The separate static sequence screen can compare a matched quasi-steady VUF only after sequence equivalents, fault impedance, measurement location and estimation window are harmonized. It cannot claim peak/transient waveform parity from one static scalar. Residuals may be nonmonotone or either sign; no direction or amount of error is predicted in advance.

ST-4 mismatch mechanisms: eliminating fast variables and using a local terminal-voltage dip omits network voltage recovery, interaction between multiple limited converters, protection and vendor active-set controls. A local pole-slip threshold and network-wide stable-recovery criterion are different observables. Current clamping can invalidate a particular small-signal reduction without proving a DAE-index change. No finite universal CCT bias is asserted.

For both tests, archive full reference traces, onset/clearing/limiter times, active modes, grounding, current priorities and post-switch regularity diagnostics if available. Evaluate step and window sensitivity and retain unsuccessful runs. Use boundary-exploration comparisons to identify where a detailed study becomes mandatory; do not score them as passing accuracy gates. All 20 matched tests remain pending until model-equivalent observables are available. An executed but mismatched local screen is INVALID for gate evaluation, not independent validation.

Harness contracts

Each case/test folder contains reduced.json and reference.json. Required fields:

```
{
  "case_id": "C0", "test_id": "ST-1",
  "operating_point_sha256": "<64-character lowercase digest>",
  "controller_sha256": "<64-character lowercase digest>",
  "disturbance_sha256": "<64-character lowercase digest>",
  "observable_definition": "mode-frequency:registered-window:v2",
  "comparison_scope": "matched-observable",
  "observable_protocol_sha256": "<64-character lowercase digest>",
  "base_mva": 100, "frequency_hz": 60,
  "solver_name": "name", "solver_version": "version",
  "solver_role": "independent-reference", "evidence_status": "executed",
  "metrics": {"mode_frequency_hz": 0.6}
}
```

The value above demonstrates syntax only and is not a supplied result. Reduced exports use solver_role=reduced-model. Other metric keys are damping_ratio, vuf_fraction, cct_s and peak_energy_gain. Hash canonical shared physical specifications; solver-specific implementations additionally need their own archived versioned source and a documented equivalence mapping. Matching hashes do not independently prove equivalence. The version-2 contract requires comparison_scope=matched-observable and a shared observable-protocol digest, covering projection, metric, measurement location, window and estimator. ST-3/4 additionally require observable_coverage={“status”:“supported”,“limitations”:“explicit model limitations”} in both records. This is a declared coverage attestation, not an automatic scientific audit. Unsupported/missing coverage or boundary-exploration scope is rejected rather than awarded a PASS. Existing version-1 result files must be reviewed and enriched; never auto-fill these claims from filenames.

Trace CSVs use time_s and explicitly named SI/per-unit channels. load_trace validates finite values, increasing time and uniform sampling. The comparator consumes registered scalar metrics; it does not automatically infer mode identity, CCT or physical energy from arbitrary CSVs.

Execution commands

```
bun publication/export-model.ts /tmp/stommel-publication
python3 publication/benchmark_ieee39.py --init /tmp/stommel-benchmark
python3 publication/benchmark_ieee39.py --compare /tmp/stommel-benchmark
python3 -m unittest discover -s publication -p 'test_*.py' -v
```

Exit codes: 0 all pairs pass; 1 at least one FAIL/INVALID; 2 missing pairs remain. Store JSON/CSV outputs without selecting only favourable cases.

Reference solver integration

PSCAD/PowerFactory require a licensed installation, a validated project deck and version-specific automation. No proprietary executable, model or credentials are available here. The portable harness deliberately accepts their exported results rather than claiming to run them. A lab must add and test its installation-specific batch exporter.

ANDES 2.0.0 now executes its bundled full IEEE39 dynamic deck: bus-14 balanced fault at 1.0 s, rf=0, xf=0.001 pu, clearing and line 14–15 trip at 1.1 s, integration through 8 s. Nominal steps 5 ms and 2.5 ms are compared on coarse timestamps. This is OPEN-RMS-1, NOT matched ST-1: allocation, controllers, equilibrium and estimation protocol differ from the reduced screen. Six positive-real eigenvalues, parameter corrections and ill-conditioning warnings remain unresolved. Completion is not stable recovery or external model validation.

NUM-C0–NUM-C3 reconstruct four reduced decks at 0%, 40%, 70% and 85% assumed inverter share in SciPy. These are not the physical C0–C3 replacement experiments above. Numerical gates in run_open_solver.py: max frequency relative discrepancy 1e-8, min simple-mode MAC 1-1e-8, max damping-formula discrepancy 1e-8, normalized eigenpair residual 1e-10 and scaled harmonic-transfer discrepancy 1e-6 (denominator max(1e-8 Hz/GW, |TypeScript gain|)). Pair transfers compare lossless dynamics; nominal IEEE39 response is withheld for indefinite symmetrized stiffness. These tolerances verify implementation, not physics. They are frozen in source before execution, but not externally timestamp-preregistered.

Run export-model.ts to produce open-solver-input.json; run run_open_solver.py –input PATH –out DIR to record numerical results, ANDES deck, matrices, spectra, raw traces and warnings. Results and step sensitivity are consolidated in Supplementary Note 12. All twenty matched ST pairs remain pending.

Reproducibility and failures

Record solver version, integration tolerance, output rate, model commit, hardware, runtime and all convergence failures. Repeat with halved steps and finer surface grids. Freeze sensitivity sweeps for inertia, load dynamics, converter gains, fault resistance, protection timing, aggregation and current priority. Provide raw traces subject to permissions. Report disagreement as evidence about the framework's limits, not a reason to retune one case. A gate-passing fit to the same data used for parameter selection is calibration, not held-out validation.

Supplementary Note 14 — Four-case selection and evidence audit

The authoritative structured ledger is `src/lib/stommel/case-framework.json`, shared by the website and conceptual figure generator. It contains source locators, mechanism chains, study questions, minimum inputs, prohibited inferences and the conventional counter-case. This is an evidence-selection framework, not a fourth executable incident model.

Why ERCOT October 2009

NERC Lesson Learned 20110705 (26 July 2011), page 1, explicitly identifies SSCI after a normally cleared fault left Type-3 wind generation radially connected through series compensation. Its Details report 485 MW across two plants, approximately 300 DFIG turbines, a single 80-mile 345-kV line with 50% series compensation, damaging oscillation growth within 200 ms and bypass 1.5 s after the initiating line trip. Page 2 describes electrical resonance, crowbar behaviour, negative damping contributed by converter controls and proposed mitigation. Primary source: https://www.nerc.com/globalassets/programs/event-analysis/lessons-learned/ll20110705_sub-synchronous_interaction.pdf . ERCOT-authored corroboration: Adams, Pappu and Dixit (2012), https://doi.org/10.1109/TDC.2012.6281678 .

No oscillation frequency is inferred from the 200-ms growth interval. The line length is geographic equipment data, not a measured electrical mode extent. The source establishes October, not an exact event day. No proprietary recording, Type-3 parameter deck, calibrated SSCI replay or measured trajectory is supplied. Independent second-reader transcription verification remains pending.

Other literature can quote a broader wind cluster, effective compensation or additional waveform estimates. Those quantities must remain attributed to their own source and system boundary; they are not substitutes for the two-plant capacity and nominal line compensation in this primary account. The selected case should not be labelled simply "weak-grid PLL instability".

Great Britain correction and overlap

ESO Technical Report Appendices, Appendix D, reproduces Ørsted's Hornsea account: insufficiently damped subsynchronous electrical resonance, depressed collection voltage and overcurrent-protection de-loading. The de-loading interval is 107 ms, approximately 5.35 cycles at 50 Hz, not subcycle. Appendix I addresses separate embedded-generation loss-of-mains responses. These observations do not make every GB event mechanism an SSCI diagnosis, nor make RoCoF Hornsea's reported de-loading trigger. Source: https://www.ofgem.gov.uk/sites/default/files/docs/eso_technical_report_-*appendices*-_final.pdf .

The case matrix deliberately acknowledges overlap: GB combines plant interaction and bulk containment; ERCOT 2009 focuses on radializing topology, series compensation, Type-3 controls and equipment damage. Four cases are contrasting study requests, not an exhaustive or mutually exclusive stability taxonomy.

Odessa model-capability correction

NERC/Texas RE, Chapter 1, documents several reduction mechanisms; PLL loss-of-synchronism trips do not alone demonstrate sustained SSCI. Chapter 2, Table 2.1, makes mechanism-specific positive-sequence capability

judgments, including a representable phase-jump-tripping entry. The main text must not claim universal RMS invisibility. EMT also needs adequate controls and protection models. Source: https://www.nerc.com/globalassets/our-work/reports/event-reports/odessa_disturbance_report.pdf . The existing replay's fault convention and prescribed loss remain unvalidated.

Selection alternatives and boundary case

Lake Bonney 2019 is not used: no plant-specific authoritative SSCI investigation was established in this audit. AEMO's documented West Murray oscillations cannot be reassigned to Lake Bonney. Xinjiang 2015 is a credible literature alternative but would require event-specific distinctions between weak-grid converter interaction, torsional consequences and hypothesized PLL contributions.

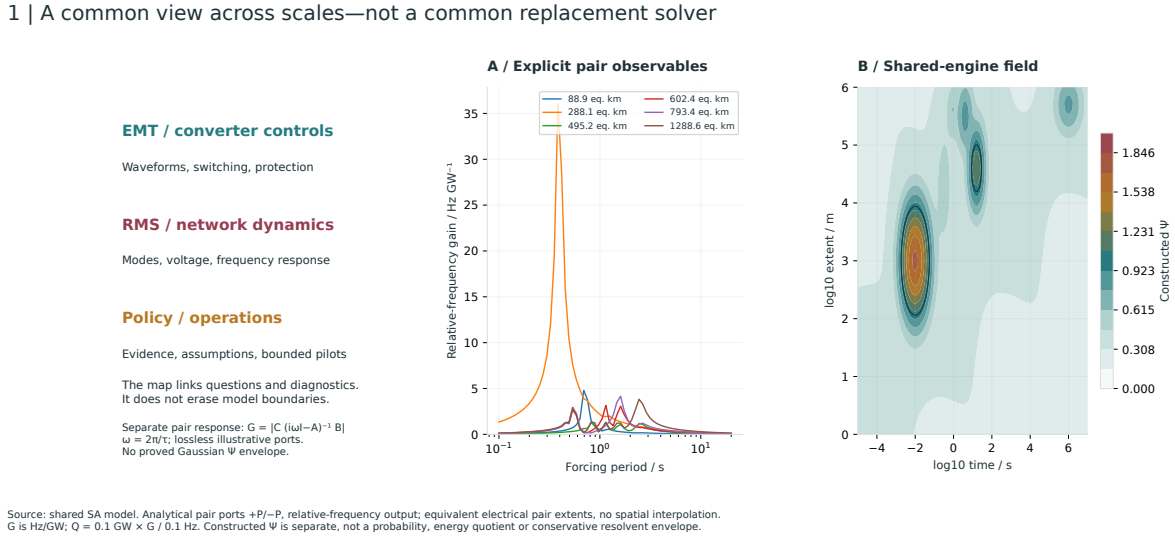
Continental Europe 8 January 2021 remains outside the four inverter/control illustrations. ENTSO-E's final investigation supplies a conventional topology/protection and synchronous-area-separation comparator: https://www.entsoe.eu/news/2021/07/15/final-report-on-the-separation-of-the-continental-europe-power-system-on-8-january-2021/ . It prevents an unsupported claim that wide-area scale transitions are uniquely converter-driven.

Evidence status and evaluation deliverables

The added source locators corroborate selected qualitative mechanisms and named timing/equipment details. They do not complete every historical numerical input audit. Retain the current replay mismatches and the 0/20 matched physical-gate status. A systematic prior-art review, independent second-reader source checks and practitioner evaluation are still required.

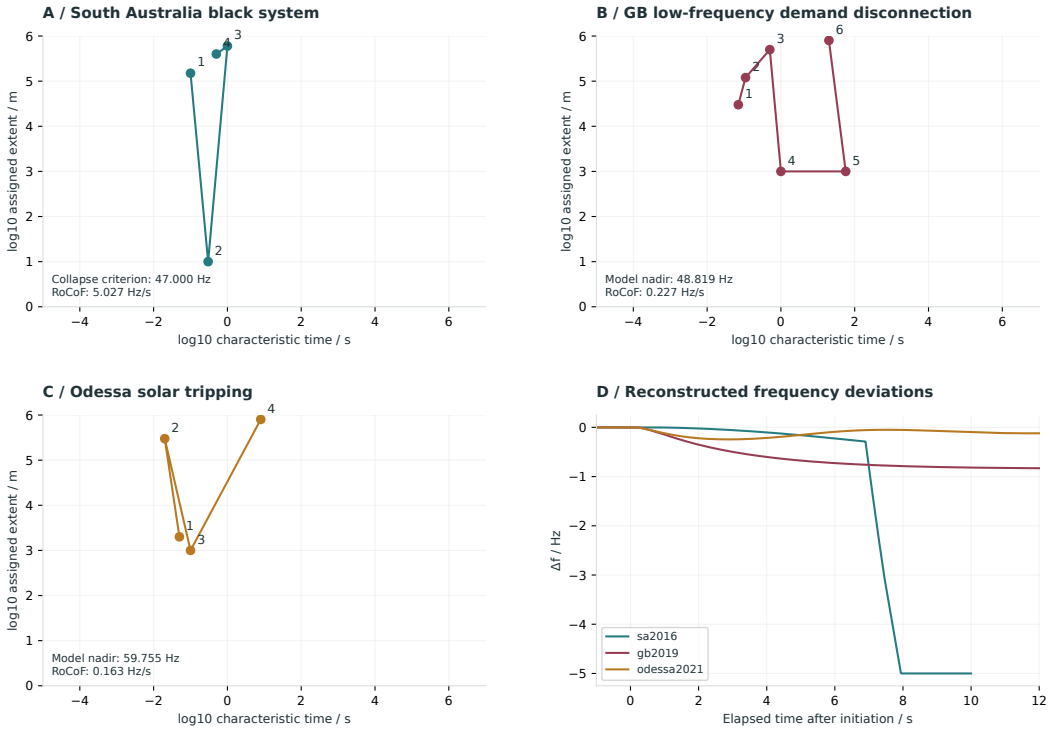
Use equivalent information in three presentation arms: report/classification/matrix, scale-and-evidence map without Ψ, and full display. Freeze task wording and acceptable study-selection answers before recruitment. Candidate tasks ask for installed repeated-fault settings (SA), plant resonance plus compound-loss containment (GB), mechanism-specific model adequacy (Odessa), radializing-outage/series-compensation controls (ERCOT) and topology/protection rather than an assumed converter cause (Europe). Score correct mechanism, appropriate study, minimum evidence and unsupported confidence separately. Pilot-informed power analysis and independent scoring precede any efficacy claim. No participant study has been executed.

Supplementary Note 15 — Technical exhibits retained outside the conceptual spine



Supplementary Figure S2 | Explicit pair observable and constructed field. Frozen model output; not a universal stability envelope. Source: unchanged V1 shared-engine snapshot.

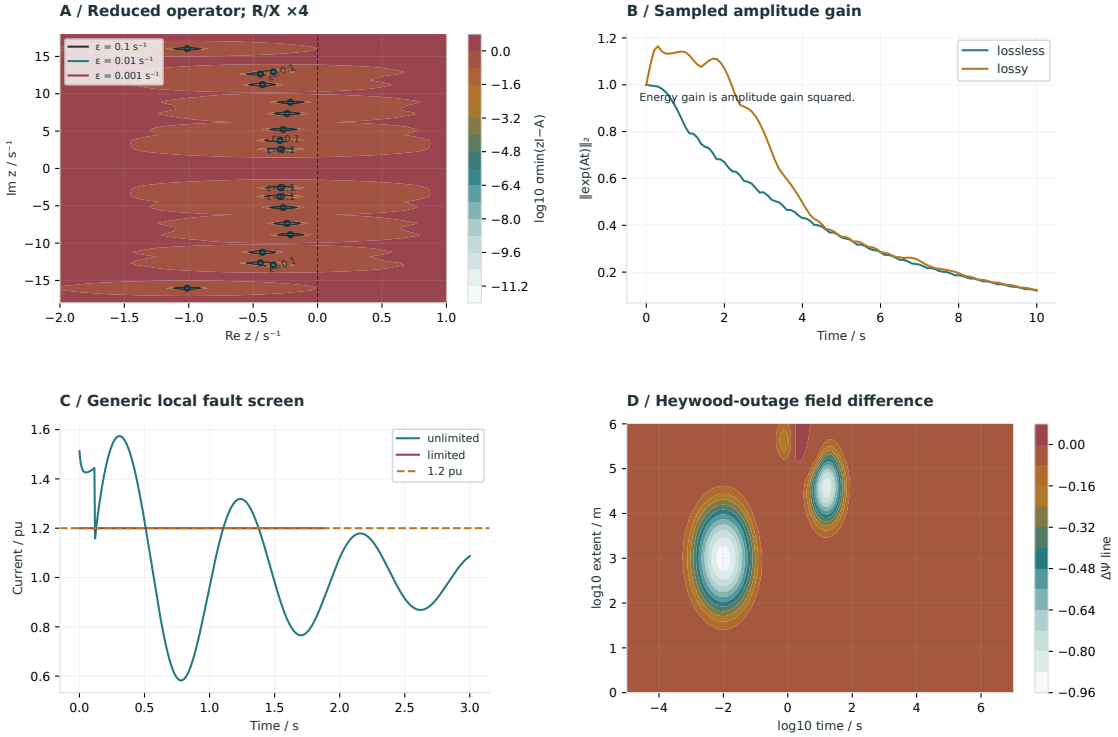
2 | Prescribed losses, assigned scales, visible model misses



Sources: official report records + current reconstruction; exact report figures await page-level audit. GB LFDD is not triggered.
SA post-collapse fall is scripted; Odessa uses a DLG convention, not a verified SLG replay. Losses are inputs, not trip predictions.

Supplementary Figure S3 | Existing prescribed incident replays. Assigned extents and scripted post-collapse decline remain explicit; GB does not trigger observed LFDD and Odessa's DLG convention is unverified. Source: unchanged V1 shared-engine snapshot; report transcriptions await complete audit.

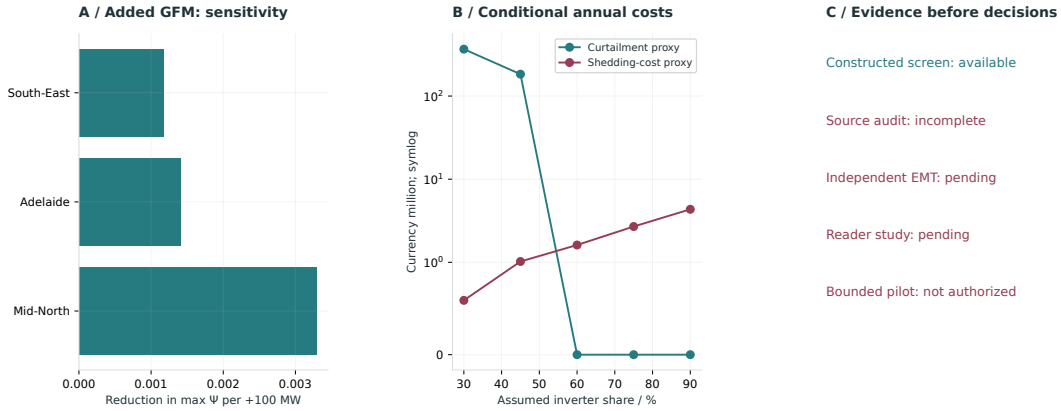
3 | Amplification and current limits: separate, conditional diagnostics



Source: generic shared model, estimated R/X conventions. Pseudospectra do not derive Ψ. Field differences are in operating conditions, not literal fractures in length–time space. No DAE-index change or full-order agreement is established.

Supplementary Figure S4 | Separate amplification and switching diagnostics. Illustrative R/X scaling and generic current limits; no DAE-index or physical fracture theorem. Source: unchanged V1 snapshot.

4 | Policy scenarios stop at the boundary of their evidence



Source: 1,000 seeded model perturbations per share; assumed VoLL 15,000/MWh, restoration 2 h, faults 0.1/year, energy value 75/MWh. Neither curve is calibrated social welfare. Nodal finite differences are not shadow prices; no queue-release estimate is supported.

Supplementary Figure S5 | Conditional policy arithmetic. Model perturbations and explicit assumed economics, not calibrated social welfare or market prices. Source: unchanged V1 snapshot.