

Compositional Lyapunov Certification of Regions of Attraction for Lossy Power Grids: A Non-Integrable Cross Term with Independent Sum-of-Squares Cross-Validation

[TODO: Author Name(s)] [TODO: Affiliation, Department, Institution, City, Country]
[TODO: corresponding author email]

Abstract—Certifying the region of attraction (ROA) of a large, lossy power grid with a monolithic Lyapunov function requires branch-and-bound (BaB) verification whose cost is exponential in state dimension. We present a *compositional* certificate, evaluated on the IEEE New England 39-bus test system Kron-reduced to its ten generator buses, that collapses an 18-D relative phase space (nine generator angles and nine speeds, one generator grounded as reference) to a 9-D, angle-only verification problem via closed-form elimination of the speed variables, making sound BaB tractable on a standard ten-generator benchmark at all. On this reduced certificate family we give a mechanistic diagnosis of why a scalar Lyapunov correction cannot compensate for lossy (non-conservative) line coupling – its required correction field is measured to be $\approx 98\%$ rotational, hence has no scalar potential – and show that moving the correction into a learned, non-integrable *vector* cross term resolves the obstruction, reaching 24.8% certified coverage of the true ROA (dense Monte Carlo ground truth on three held-out angle slices) at 100% soundness, within 0.5 percentage points of that certificate family’s own zero-relaxation verification ceiling. We cross-validate the entire pipeline with an independently implemented, exact sum-of-squares (Positivstellensatz) certificate on the same polynomial, which *ties* the interval BaB result on all three evaluation slices individually (mean coverage 22.0% vs. 22.0%; e.g. 93.8° vs. 93.6° proved radius on one slice) – to our knowledge the first exact, non-interval-relaxed cross-check in this line of work. Finally, we show certified coverage decomposes into two *independent* sub-problems – how far stability can be proved (radius) and how well the proved region’s geometry matches the true boundary (shape) – and that a single convex quadratic program, run on all three slices, recovers a sound 86% relative coverage gain (mean 39.8% vs. 21.4%) on the shape sub-problem alone, a gain invisible to six independent gradient-based searches that optimized both sub-problems jointly. Both the physics-only and the learned-correction mechanisms transfer to a second, independently constructed benchmark topology (the IEEE/WSCC 9-bus system), and the trained certificate is robust to a $\pm 15\%$ loading perturbation on two of three evaluation slices, collapsing on the third – reported exactly rather than smoothed into a single robustness claim. We report this decomposition, together with eight independently designed negative results on the radius sub-problem, as the paper’s central structural contribution, and outline what would be required to close the remaining gap to the true ROA boundary.

Index Terms—Lyapunov methods, power system transient sta-

bility, region of attraction, sum-of-squares programming, formal verification, neural certificates, branch-and-bound, semidefinite programming.

I. INTRODUCTION

TRANSIENT stability assessment asks whether, after a disturbance, a power grid returns to synchronous operation, and from how large a set of post-fault states. Two classical tools answer parts of this question. Time-domain simulation is fast and exact for the states it simulates, but certifies nothing about states it did not. The transient energy function (TEF) method [1] gives an analytic Lyapunov function and a provably correct stability boundary – but only for *lossless* networks, where the coupling admits an exact scalar potential. Real transmission lines have nonzero conductance, and no exact energy function exists once loss is present; this is the starting obstacle this paper addresses directly.

More recent work offers two paths around the lossless restriction. Neural Lyapunov / neural certificate methods [2], [3] learn a candidate certificate and verify $\dot{V} < 0$ with sound bound propagation – interval arithmetic, or tighter linear-relaxation methods such as CROWN [4] and its branch-and-bound extensions built on the Reluplex line of work [5], [6] – but verification cost is exponential in state dimension, which in practice caps these methods at a handful of generators. Sum-of-squares (SOS) and occupation-measure methods [7], [8] compute provably convergent *outer* approximations of the ROA via semidefinite programming (SDP), with no interval relaxation gap at all, but the moment matrix grows combinatorially in state dimension and polynomial degree, which caps them earlier still. Neither line has, to our knowledge, been applied to a standard ten-generator *lossy* benchmark – the IEEE New England 39-bus system [9] – with a rigorous, sound accounting of how much of the true ROA the resulting certificate actually captures.

This paper makes four contributions.

- 1) A *compositional reduction*: for any Lyapunov function quadratic in the speed variables, the speed dimension eliminates in closed form (Section III), collapsing an 18-D verification problem to 9-D and making sound BaB tractable on this system size at all.
- 2) A mechanistic diagnosis of why a scalar Lyapunov correction cannot fix a lossy grid (Section III-B), and

a working fix – a non-integrable vector cross term – reaching 24.8% certified coverage at 100% soundness, within 0.5 points of the certificate family’s own zero-relaxation ceiling (Section IV).

- 3) An independently implemented, exact SOS/Positivstellensatz certificate that cross-validates the interval BaB result to within measurement precision (Section V) – the first exact-method confirmation in this line of work, and evidence the 93–106° band is a property of the certificate rather than an artifact of either verifier.
- 4) A structural decomposition of certified coverage into an independent *radius* question and *shape* question (Section VI), with a single convex program, run on all three evaluation slices, recovering a sound 86% relative coverage gain on the shape question alone – a gain six independent gradient-based searches missed because they optimized both questions jointly (Section VII).

II. SYSTEM MODEL AND PROBLEM STATEMENT

The plant is the swing-equation model of the **IEEE New England 39-bus test system** [9], Kron-reduced to its ten generator buses [10], with a closed-loop droop equilibrium and controller gains reused verbatim from Cui’s Lyapunov-regularized reinforcement learning benchmark [TODO: verify exact citation for Cui’s Lyapunov-RL paper] – a standard transient-stability test case, not a bespoke or anonymous network:

$$M_i \ddot{\delta}_i = P_i(\delta) - d_i \dot{\delta}_i, \quad (1)$$

$$P_i(\delta) = P_{m,i} - \sum_j [F_{ij} \sin \delta_{ij} + G_{ij} \cos \delta_{ij}] \quad (2)$$

with $\delta_{ij} := \delta_i - \delta_j$, F_{ij} the (conservative) susceptance coupling, and G_{ij} the transfer *conductance* – identically zero only in the idealized lossless limit. One generator is grounded as an angle reference, leaving a 9-D relative angle vector ϕ and 9-D relative speed vector ω as the active state; we write $\delta = \delta^* + \phi$ throughout, with δ^* a fixed equilibrium.

Region of attraction. For the state $x = (\phi, \omega)$, the ROA is $\{x_0 : \lim_{t \rightarrow \infty} x(t; x_0) = 0\}$ under (1). We estimate its true extent by dense Monte Carlo integration (12 s horizon, fourth-order Runge–Kutta, $\Delta t = 0.02$ s) from angle-only initial conditions ($\omega_0 = 0$) on three held-out 2-D angle slices – the generator pairs (5, 6), (4, 7), (1, 9) in the active index set – with the seven remaining active angles pinned at δ^* . A grid point is labeled stable if the trajectory returns within 0.5 rad angular error and 0.05 rad/s speed error of equilibrium.

Coverage and soundness. For a candidate certified set $\mathcal{C} \subseteq \{\phi\}$, *coverage* is $|\mathcal{C} \cap \text{stable}|/|\text{stable}|$ on the sampled grid, averaged over the three slices; *soundness* is $100\% - 100\% \cdot |\mathcal{C} \cap \text{unstable}|/|\mathcal{C}|$, and every number reported as “100% sound” in this paper has zero sampled points in $\mathcal{C} \cap \text{unstable}$.

III. COMPOSITIONAL CERTIFICATE STRUCTURE

A. Closed-form speed elimination

Consider the certificate family

$$V(\phi, \omega) = \frac{1}{2} \omega^\top M \omega + U(\phi) + \Psi(\phi)^\top M \omega, \quad (3)$$

with $M = \text{diag}(M_i)$, U a scalar potential, and Ψ a vector field – the *cross term*. Because (3) is quadratic in ω by construction, differentiating along (1) gives

$$\dot{V}(\phi, \omega) = \omega^\top Q(\phi) \omega + b(\phi)^\top \omega + c_0(\phi), \quad (4)$$

$$J := \partial \Psi / \partial \phi, \quad (5)$$

$$Q = -\text{diag}(d) + \Omega \text{sym}(MJ), \quad (6)$$

$$b_i = -L_i(\phi) - d_i \Psi_i(\phi), \quad (7)$$

$$L_i(\phi) := \sum_j G_{ij} (\cos \delta_{ij} - \cos \delta_{ij}^*), \quad (8)$$

$$c_0 = \Psi(\phi)^\top P(\phi), \quad (9)$$

where Ω is the per-unit frequency-base constant relating $\dot{\phi} = \Omega \omega$. For $Q \prec 0$, (4) is a concave quadratic in ω and its maximum over *every* speed has closed form:

$$h(\phi) := \max_{\omega} \dot{V} = \frac{1}{4} b(\phi)^\top (-Q(\phi))^{-1} b(\phi) + c_0(\phi). \quad (10)$$

Proving $h(\phi) < 0$ over an angle-only box proves $\dot{V} < 0$ for *every* speed in that box at every angle in it: the 9-D speed dimension is eliminated analytically, not searched. This is the single structural choice that makes sound BaB tractable on ten generators; the same claim verified directly over the full 18-D state (Section VII, last row) is the exponential regime every prior neural-Lyapunov result on this system size runs into.

B. The curl obstruction

Take $\Psi \equiv 0$ and U the classical (lossless) TEF. The lossy residual $L_i(\phi)$ above enters b unmodified and h *squared and positive*, so it can only raise the certified threshold’s cost, never lower it. A scalar residual $\Phi(\phi)$ added to U contributes only a *gradient* $\Omega \nabla \Phi$ to b ; the field required to cancel L has curl

$$\nabla \times F_{\text{req}} = 2 G_{ij} \sin \delta_{ij}, \quad (11)$$

exactly the lossy conductance, measured at $\approx 98\%$ of the field’s magnitude on this system (median ratio over the evaluation slices). No gradient supplies a field that is almost entirely rotational: a scalar correction is provably the wrong object, independent of training procedure or capacity. We verified this is not an optimization artifact by isolating the term (removing the coverage objective and running 700 steps of pure adversarial h -minimization with full network capacity): the residual converges to strictly *worse* than the uncorrected physics baseline (9.4% vs. 14.6% coverage; Table VI, row c8).

C. Non-integrable vector cross term

Moving the correction into Ψ removes the integrability requirement: $b_i = -L_i - d_i \Psi_i(\phi)$ is a *pointwise algebraic* condition on Ψ_i ’s value, not a condition that some field be a gradient. We parametrize Ψ_i as a shallow (one hidden layer, tanh) network per generator over local port variables (own angle plus nearest electrical neighbours), initialized at the affine physics baseline $\Psi_i = \epsilon \phi_i$ so training can only bank improvements, and trained counterexample-guided against h evaluated through the network.

TABLE I
MONOLITHIC BRANCH-AND-BOUND LEAF-BOX COUNT VS.
COMPOSITIONAL OPERATION COUNT, SAME VERIFICATION CLAIM,
MEASURED ON THIS SYSTEM.

Dim.	Active gens.	Monolithic leaves	Speedup
2	1	64	1×
6	3	2.6×10^5	$1.2 \times 10^3 \times$
10	5	1.1×10^9	$2.4 \times 10^6 \times$
14	7	6.9×10^{10}	$1.1 \times 10^8 \times$
20	10	1.15×10^{18}	$7.0 \times 10^{14} \times$

D. Sound interval verification

$h(\phi)$ is verified by interval branch-and-bound: $(-Q)^{-1}$ is never interval-inverted directly (numerically loose); instead $-Q = D - E$ is split into its diagonal and off-diagonal parts and E is dominated by a per-row Gershgorin bound in Loewner order, giving a sound closed-form upper bound on h per box. A box is proved once its upper bound is negative; boxes that fail are adaptively bisected. A small central cube around the equilibrium ($|\phi_i| \leq \rho = 0.05$) is excluded from BaB – $h \rightarrow 0$ there by construction, so no box can show strict negativity – and covered separately by a local Hessian argument plus dense sampling, consistent across every certificate in this paper.

E. Why compositional, not monolithic: a measured cost comparison

The reduction in Section III is not a convenience – it is what makes verification possible on this system size at all. We measured this directly rather than asserting it asymptotically: a naive worst-case small-gain composition (bounding each coupling term by its magnitude, independent of the physical structure) reproduces the same verification claim at increasing generator count on this exact system.

At 20-D, monolithic branch-and-bound needs on the order of 10^{18} leaf boxes to resolve the same claim the compositional reduction resolves in $O(10^3)$ operations – a measured, not asymptotic, $7 \times 10^{14} \times$ gap. *This scaling mechanism is not itself sufficient for a good certificate*: the same measurement shows worst-case small-gain composition degrades quickly (certifying only through 4-D before the composed bound goes positive), which is exactly why Sections III-B and III-C replace worst-case composition with the physics- structured, then learned, cross term. The scaling argument establishes why monolithic verification is not an option here; it does not by itself establish that the resulting certificate is a good one.

An empirical, not asymptotic, data point on the identical system. A companion verification effort in the accompanying repository¹ independently trained and *monolithically* verified a neural Lyapunov function on this same network using CROWN / α, β -CROWN linear-relaxation bound propagation (auto_LiRPA) [4], [6] – a materially different, and generally tighter, bounding technology than the interval arithmetic used

throughout this paper. It certified cleanly at one generator (4-D full state: one generator’s angle and speed) and, without degradation, at two generators (also 4-D $\rightarrow$ 4-D in their own reduced coordinates, i.e. the next rung up), with branch-and-bound subdomain count growing $\approx 5 \times$ (roughly 10 to 35–61 across five seeds) and wall-clock per seed growing from ≈ 6 to ≈ 10 minutes – real, measurable growth, but not yet an explosion. Extending to three generators was explicitly deferred as the next rung and was not attempted, so this is not a proven wall; it is the honest extent of what monolithic CROWN verification has actually been run to on this system. The compositional reduction in this paper reaches all nine active generators.

IV. RESULTS I: CERTIFIED COVERAGE

Constraining the optimizer to *retain* full structural (coupling LMI) margins improved coverage rather than costing it (Table II, last row vs. second row) – the unconstrained run had been spending margin it did not need to spend.

A. Verifier-tightness ceiling

Before searching further, we measured the largest radius *any* verifier – interval, exact, or oracle – could ever certify for a given field, by dense-grid evaluation of the exact (unrelaxed) condition $h(\phi) < 0$. This zero-relaxation ceiling upper-bounds every possible verifier for that field.

A perfect, zero-relaxation verifier buys at most half a coverage point here. **The cap is the certificate’s own sublevel set, not the bounding.**

V. INDEPENDENT VERIFICATION: SUM-OF-SQUARES CROSS-VALIDATION

Every result to this point uses one verification technology (interval arithmetic) and six architecturally distinct attempts to beat the 24.8% ceiling (Section VII), all converging on the same band. That convergence is suggestive but not dispositive: every one of those attempts still used interval arithmetic somewhere in the loop.

We therefore built an independent, exact **Positivstellensatz certificate**. Restricting to a single 2-D slice at a time and lifting the two varying angles to $(\sin, \cos)$ pairs (s_1, c_1, s_2, c_2) via product-to-sum identities, and replacing the (non-polynomial) affine cross term with $\Psi_i = \epsilon \sin \phi_i$ (bounded and low-degree once lifted), $h(\phi)$ restricted to a slice is an exact polynomial of total degree 5. A degree-6 sum-of-squares program then searches for multipliers $\sigma_0, \sigma_1, \sigma_2, \lambda_1, \lambda_2$ and a slack γ such that

$$-H - \gamma = \sigma_0 + \sigma_1 g_1 + \sigma_2 g_2 + \lambda_1 h_1 + \lambda_2 h_2, \quad (12)$$

$\sigma_k \succeq 0$ (SOS), $H := 4q_1 q_2 h$ the denominator-cleared polynomial, g_k the box/annulus inequalities, and h_k the unit-circle identities $s_k^2 + c_k^2 = 1$; γ is *maximized* rather than fixed, which is both a genuine safety margin and numerically far more robust for interior-point solvers than a pure feasibility query. Following the interval verifier’s own convention, the central neighbourhood is excluded via two region certificates (one

¹Not part of this paper’s own results; logged for context because it uses the identical IEEE 39-bus benchmark and provides a genuine head-to-head reference point, not a proxy.

TABLE II
HELD-OUT PROVED RADII AND MEAN COVERAGE, PHYSICS-ONLY FIELD VS. TRAINED VECTOR CROSS TERM.

Field	LMI margin	(5, 6)	(4, 7)	(1, 9)	Coverage
Physics only	100% / 100%	102.0°	96.4°	75.1°	23.2%
Learned, unconstrained	78% / 66%	105.3°	100.8°	69.5°	24.0%
Learned, floor-constrained	100 / 100%	106.4°	103.1°	72.9°	24.8%

TABLE III
ZERO-RELAXATION VERIFIER CEILING VS. PROVED COVERAGE.

Field	Ceiling	Proved	Headroom
Physics (exactly affine)	23.3%	23.2%	0.1 pts
Best learned field	25.3%	24.8%	0.5 pts

TABLE IV
INTERVAL BAB VS. INDEPENDENT SOS CERTIFICATE, IDENTICAL POLYNOMIAL ($\Psi_i = \epsilon \sin \phi_i$), ALL THREE EVALUATION SLICES.

Slice	BaB radius	BaB cov.	SOS radius	SOS cov.
(5, 6)	93.6°	27.8%	93.8°	28.0%
(4, 7)	88.9°	20.4%	88.5°	20.2%
(1, 9)	85.6°	17.7%	85.6°	17.7%
Mean		22.0%		22.0%

per axis) whose union covers the annulus, matching BaB’s box exclusion exactly rather than the (unprovable) full disk including the equilibrium. We run this on all three evaluation slices, not a single pilot.

The two methods tie on every individual slice, not merely in aggregate – the strongest form of the cross-validation result – and agree to within bisection precision throughout. (These coverage figures use the SOS-compatible $\epsilon \sin \phi_i$ baseline rather than the $\epsilon \phi_i$ baseline behind Table II’s 24.8% headline, so the two are not directly comparable in absolute terms; this table is a controlled verifier-vs-verifier comparison on one fixed object, not a new coverage claim.) This is, to our knowledge, the first exact (non-interval) confirmation that a BaB-proved radius on this system is not leaving verifiable region on the table – evidence the 93–106° band (Table II) is a property of the certificate rather than an artifact of branch-and-bound’s own conservatism.

VI. SHAPE AND RADIUS ARE INDEPENDENT QUESTIONS

Coverage has depended on two logically separate quantities throughout: how far $h < 0$ can be proved (the *radius*, Sections IV and V), and how well the certified sublevel set’s *geometry* matches the true stable region at *that radius* (the *shape*). Crucially, coverage is defined via $U(\phi)$ ’s sublevel set alone – the cross term Ψ never enters that definition. Six gradient-based attempts to reshape U (Section VII, rows c15/c15b) found nothing distinguishable from noise, because gradient descent was solving both questions in one entangled, non-convex objective.

Decoupling them: fix the radius at each slice’s own already-proved boundary r^* (Table IV, minus a 0.02 rad safety mar-

TABLE V
COVERAGE AT EACH SLICE’S OWN FIXED PROVED RADIUS r^* .

Slice	r^*	Classical TEF	Naive disk	SOS-fit shape
(5, 6)	92.5°	27.1%	41.6%	41.6%
(4, 7)	87.4°	19.8%	40.5%	40.5%
(1, 9)	84.5°	17.4%	37.3%	37.3%
Mean		21.4%	39.8%	39.8%

gin), then pose the shape question alone as a single convex program,

$$p(\phi) = z(\phi)^\top Q z(\phi), \quad Q \succeq 0, \quad (13)$$

with z the degree- ≤ 2 monomial basis in (s_1, c_1, s_2, c_2) (a degree-4 polynomial with provably convex sublevel sets), fit to Monte Carlo ground truth by minimizing hinge loss on stable points subject to a hard margin against every sampled unstable point and a dense-grid containment constraint keeping $\{p \leq 1\}$ inside r^* . Run on all three slices:

An **86% relative gain** over the classical shape (mean 39.8% vs. 21.4%), fully sound on all three slices (zero certified-unstable points throughout) and found by a single convex solve per slice. Containment (sampling-based, not itself a Positivstellensatz certificate) held cleanly on two of three slices; slice (5, 6) showed a borderline leak of 0.006 rad ($\approx 0.34^\circ$) beyond r^* on the dense post-hoc check, smaller than that check grid’s own spacing and plausibly a discretization artifact rather than a genuine violation – reported exactly rather than rounded away, consistent with this paper’s practice throughout.

Reported honestly rather than rounded away: at every proved radius available in this study ($\leq 106.4^\circ$), the sampled ground truth inside each disk is *entirely stable* – the nearest actual unstable point on slice (5, 6) lies at 114.9° , beyond any radius proved by any method in this paper. The fitted shape therefore matches the naive disk almost exactly on every slice, with no boundary information yet available to differentiate a good shape from a merely large one. Table V is a real, sound statement about how badly the classical shape under-uses a given proved radius; it is not yet evidence that shape optimization alone can extend the radius toward the true boundary.

VII. NEGATIVE-RESULT LEDGER

We report eight independently designed attempts to extend the radius beyond Table II’s 24.8%, all negative, as findings rather than discarded drafts – their convergence is itself evidence about where the certificate family’s limit lies.

Row c17 is the most structurally permissive variation tested: it removes the very assumption (quadratic-in- ω) that Section III relies on for tractability, verifying $\dot{V} < 0$ by direct

TABLE VI
FULL EXPERIMENTAL LEDGER. ALL COVERAGE FIGURES ARE 100% SOUND UNLESS NOTED.

ID	Tested	Result	Outcome
c1–c6	Compositional scaling ladder: small-gain $\rightarrow$ conservative-core $\rightarrow$ structured Lyapunov proof $\rightarrow$ exact energy function $\rightarrow$ certified ROA $\rightarrow$ dissipativity composition	10-D reached, physics-structured proof (no SDP)	Foundation
c7	Coverage vs. dense Monte Carlo ground truth, corrected adversarial threshold	14.6% physics baseline (prior 22.4% retracted – measured at an invalid threshold)	Baseline
c8	Scalar potential residual $\Phi(\phi)$	Strictly worse in isolation (9.4%); curl obstruction identified	Ruled out
c9–c11	Non-integrable vector cross term, verified in-network by interval BaB, margin-floor training	24.8%, 100% sound	Headline
c12	Zero-relaxation verifier ceiling	0.5-pt headroom only	Diagnosis
c13	Stress test: margin floor above 100%, deep 9-D BaB depth probe	Margin ceiling 100–110%; 9-D wall is combinatorial (523,775 boxes at 5.7°)	Diagnosis
c14	Exact closed-form cross-term cancellation, zero parameters	$b \equiv 0$ (to 10^{-15}) but c_0 term explodes; ceiling $\rightarrow$ 0.1%	Ruled out
c14b	Scalar blend of exact term with physics baseline	Flat, 23.6% peak vs. 23.3% baseline	Ruled out
c15/c15b	Potential-reshaping residual, frozen and jointly-trained cross term	Flat at 16.8%; best checkpoint = untrained start	Ruled out
c16	Remove per-node locality; full-state visibility, $2\times$ capacity	17.4% adversarial, training plateaued at step 150/900	Ruled out
c17	Remove the quadratic-in- ω restriction entirely; free $V(\phi, \omega)$, full 18-D, exact autograd $\dot{V}$	Identical to baseline to one decimal (15.7%)	Ruled out
c18	Independent SOS/Positivstellensatz certificate vs. interval BaB, all 3 slices	tie on every slice, mean 22.0% vs. 22.0%	Cross-validated
c19	Convex quadratic-form shape search at each slice’s own proved radius, all 3 slices	+86% relative (39.8% vs. 21.4%), sound	New lever

TABLE VII
OPERATING-POINT ROBUSTNESS, TRAINED FIELD, NO RETRAINING.

Spread	(5, 6)	(4, 7)	(1, 9)	Mean
$\times 0.85$	106.4°	105.3°	0.0°	70.6° (75%)
$\times 1.00$	106.4°	103.1°	72.9°	94.1° (100%)
$\times 1.15$	105.3°	100.8°	0.0°	68.7° (73%)

automatic differentiation against the true nonlinear dynamics over the full, correctly-projected joint sublevel set $\{V(\phi, \omega) \leq c\}$ (an earlier, incorrect version that boxed ω independently of ϕ produced a spurious near-zero result, traced to checking physically irrelevant high-kinetic-energy points; corrected before the result in Table VI was trusted). That even this maximally flexible certificate, with a general neural correction over the full state, matches the fixed-structure baseline to one decimal place is the strongest single piece of evidence that the ceiling is a property of the underlying verification problem on this system, not of any one certificate architecture.

VIII. GENERALIZATION: OPERATING-POINT ROBUSTNESS AND A SECOND TOPOLOGY

A. Operating-point robustness of the trained field

We reused the operating-point robustness method established for an earlier, simpler certificate in this project’s own development (scale the equilibrium angle deviations from their mean by a factor sp , rebuild the linearized coupling, re-certify – this avoids the droop-only steady-state trap, since a raw loading change under pure droop control has no exact fixed point to linearize at) and applied it, for the first time, to the actual headline result: c11’s trained vector cross term, with no retraining.

Reported exactly, not smoothed into a single robustness claim: two of three slices are essentially fully robust across a $\pm 15\%$ loading band (under 5% change from nominal). The third – slice (1, 9), already the hardest slice throughout this study – **collapses to 0° under either perturbation direction tested**. The trained field is robust where it is robust and locally overfit to the exact nominal equilibrium where it is not; both are reported, not just the former.

B. A second, independent benchmark topology

Every result to this point is on one network. We built a second, fully independent classical multimachine model – the IEEE/WSCC 9-bus 3-generator system [11] – from verified external sources: network topology, impedances, and generator setpoints from MATPOWER’s official `case9.m`; generator inertia and transient reactance cross-confirmed from two independent published sources; the AC power flow solved by `pandapower` (a tested solver, not a hand derivation); the classical Kron reduction to generator-internal EMF nodes performed with two verification checks before anything downstream was trusted: an exact equilibrium-balance check (residual 10^{-16} after a units bug was caught and fixed – an intermediate version conflated terminal power with internal-node power balance, caught because the residual matched each generator’s self-admittance exactly) and a Hurwitz stability check on the reference-grounded system (an intermediate version showed exactly-zero eigenvalues at every damping level tested, traced to never having grounded a reference generator – the same rotational-symmetry zero mode this project independently discovered for the 39-bus system, not a new bug).

With only two active generators after reference-grounding, the full relative state is the certified 2-D space – no slicing required.

TABLE VIII
SECOND TOPOLOGY (IEEE/WSCC 9-BUS), PHYSICS-ONLY VS. LEARNED
CROSS TERM.

Field	Radius	Coverage	Sound
Physics only	88.5°	23.4%	100%
Learned cross term	104.7°	32.3%	100%

The physics-only baseline (23.4%) lands remarkably close to the 39-bus system’s own physics baseline (23.2%) despite the two networks sharing no topology, size, or parameters. The learned cross term – the identical training procedure from Section III-C, reused unchanged – improves this to 32.3%, a **38% relative gain**, larger than the 39-bus system’s own 7% relative headline gain. Both the physics-only mechanism and the learned-correction mechanism transfer to an independently constructed topology. Scope, stated plainly: only the physics-baseline and learned-cross-term stages were re-run here, not the full SOS/shape-search arc, and the coupling-LMI margin-retention check (Section IV) was not re-verified for this system, since that specific machinery is tied to the 39-bus data format – offered as evidence the reduction and verification *approach* generalizes, not as a second complete experimental arc.

IX. DISCUSSION AND LIMITATIONS

- **Two topologies, and operating-point robustness on the primary one – not an exhaustive sweep.** Section VIII adds a second, independent network and a loading-condition band, but this remains short of a systematic study across many topologies and conditions.
- **Coverage is a 2-D-slice metric.** All coverage figures are fractions of Monte-Carlo-sampled points on three 2-D angle slices, not a volume estimate of the full 9-D (or 18-D) state space; raw high-dimensional volume is not a well-defined target without a reference measure, and every number in this paper is reported on the same slices for comparability.
- **The central neighbourhood is not BaB- or SOS-proved.** Both verifiers exclude a small region around the equilibrium (where the margin vanishes by construction) and rely on a local-Hessian-plus-dense-sampling argument there, consistently across every certificate reported.
- **The row-c17 speed bound is empirical.** The full-state certificate is verified over a speed range set from simulation (peak $|\omega| = 1.50$ observed over 200 trajectories, verified to 2.0), not from a forward-invariance proof.
- **The Section VI containment is sampling-based.** The fitted shape’s containment inside r^* is checked on a dense grid, not by a Positivstellensatz containment certificate – the one place in the shape result not yet fully rigorous.
- **Eight negative results bound one certificate family, not all certificates.** They constrain “transient-energy potential plus some correction, quadratic in ω ” specifically (row c17 included, since it still uses the same V_{pot} backbone); a certificate abandoning that structure entirely – a different potential backbone verified end-to-end, a

max-of-quadratics or multiple-Lyapunov-function construction, or a non-Lyapunov paradigm such as Hamilton–Jacobi reachability or occupation-measure ROA estimation – is not excluded by anything measured here, and would face its own dimensionality limits at 9–18 D.

X. CONCLUSION

The radius question is now unusually well bounded: eight independent methods, including one exact verification technology (Section V) and the single most structurally permissive change available inside the Lyapunov paradigm (row c17), converge on 93–106°. The shape question, in contrast, is genuinely open: the convex program in Section VI is a first, minimal instance of a much larger family – higher-degree SOS shapes, or a proper bilinear alternation between shape and cross term – and has not yet been tested at a radius where real boundary information exists to fit against. Extending the proved radius even modestly and re-running the same convex shape fit there is the most direct next step this study’s own evidence points to.

DATA AND CODE AVAILABILITY

[TODO: repository URL and license]. All experiment scripts, raw results (JSON), and the running research log referenced throughout this paper are version-controlled in the accompanying repository.

ACKNOWLEDGMENT

[TODO: funding / acknowledgments].

REFERENCES

- [1] H.-D. Chiang, C.-C. Chu, and G. Cauley, “Direct stability analysis of electric power systems using energy functions: theory, applications, and perspective,” *Proc. IEEE*, vol. 83, no. 11, pp. 1497–1529, 1995.
- [2] Y.-C. Chang, N. Roohi, and S. Gao, “Neural Lyapunov control,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [3] C. Dawson, S. Gao, and C. Fan, “Safe control with learned certificates: a survey of neural Lyapunov, barrier, and contraction methods for robotics and control,” *IEEE Trans. Robotics*, 2023.
- [4] H. Zhang, T.-W. Weng, P.-Y. Chen, C.-J. Hsieh, and L. Daniel, “Efficient neural network robustness certification with general activation functions,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [5] G. Katz, C. Barrett, D. L. Dill, K. Julian, and M. J. Kochenderfer, “Reluplex: an efficient SMT solver for verifying deep neural networks,” in *Computer Aided Verification (CAV)*, 2017.
- [6] K. Xu, Z. Shi, H. Zhang, Y. Wang, K.-W. Chang, M. Huang, B. Kailkhura, X. Lin, and C.-J. Hsieh, “Automatic perturbation analysis for scalable certified robustness and beyond,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [7] A. Papachristodoulou and S. Prajna, “On the construction of Lyapunov functions using the sum of squares decomposition,” in *IEEE Conference on Decision and Control (CDC)*, 2002.
- [8] D. Henrion and M. Korda, “Convex computation of the region of attraction of polynomial control systems,” *IEEE Trans. Automatic Control*, vol. 59, no. 2, pp. 297–312, 2014.
- [9] T. Athay, R. Podmore, and S. Virmani, “A practical method for the direct analysis of transient stability,” *IEEE Trans. Power Apparatus and Systems*, vol. PAS-98, no. 2, pp. 573–584, 1979. [TODO: verify page range before submission] – origin of the IEEE New England 39-bus test system.
- [10] J. H. Chow, *Power System Toolbox*. [TODO: verify citation form – software/toolbox documentation, Rensselaer Polytechnic Institute / Cherry Tree Scientific Software] – source of the NE39 network data used in the Kron reduction to ten generator buses.

- [11] P. M. Anderson and A. A. Fouad, *Power System Control and Stability*, 2nd ed. Piscataway, NJ: IEEE Press, 2003. – origin of the IEEE/WSCC 9-bus 3-machine test system; network data cross-checked against MATPOWER's `case9.m` (github.com/MATPOWER/matpower).