STABILITY OF ALMOST-FLAT LOSSY COCYCLES: AN INDEX OBSTRUCTION SURVIVES RANK LOSS

JUAN CARLOS PAREDES

ABSTRACT. The stability theory of almost-representations — when is a map that almost satisfies a group's relations close to one that satisfies them exactly? — has reached a sharp modern form: for higher-rank lattices, pointwise Frobenius-norm stability follows from second-cohomology vanishing, pointwise operator-norm stability fails through index obstructions, and uniform stability holds for all submultiplicative norms. Every result in this landscape assumes invertible targets: unitaries. We formulate the corresponding stability problem for *lossy* targets — rank-deficient partial isometries, the maps that arise when translations between heterogeneous systems destroy information — over graphs with cycle relations. A structure lemma identifies the exact objects: a family of lossy edge maps is exactly consistent if and only if it has the form T_e = V_j V_i$^\mathsf{T}$ for orthonormal frames V at the nodes, in which case every cycle composite is a projection. Against that family we measure stability across the grid of (matrix norm) $\times$ (defect functional), with planted-obstruction, noise-floor, and invertible-column controls. Dense systems are uniformly stable with a bounded penalty for rank deficiency. On the torus with a unit-Chern flux planted in the surviving subspace, the local defect vanishes with system size while the operator-norm distance to the nearest exact family stays pinned at order one — identically at full rank, corank one, and half rank — while the Frobenius distance decays. We conjecture that the pointwise Frobenius/operator dichotomy of the invertible theory, including its index mechanism, survives rank loss. The conjecture is machine-checked down to a single estimate: a Lean 4 formalization proves that it follows — uniformly in system size, with the constant preserved — from one transported rigidity statement (a dimension-free lower bound on the distance from the flux magnetic-translation pair to every commuting contraction pair), which direct adversarial optimization over commuting contraction pairs fails to falsify: no competitor below the trivial zero-pair distance of exactly 1 was found at any size or rank-padding, so the sought constant lies in (0, 1]. The open problem this paper poses to the stability community is exactly that estimate.

1. INTRODUCTION

Call a map from a finitely presented group to a metric group an *almost-representation* if it nearly satisfies the defining relations, and call the group *stable* if every almost-representation is a small deformation of an exact one. Which groups are stable turns out to depend — remarkably — not only on the group but on how "nearly" and "small" are measured. For higher-rank lattices the modern landscape is sharp: vanishing of second cohomology with unitary coefficients implies pointwise stability in the Frobenius norm [DCGLT]; nonvanishing even-degree cohomology produces pointwise *instability* in the operator norm [BLSW; Dadarlat], with the obstruction ultimately of index type [Voiculescu; Exel–Loring]; and uniform stability — small defect over the whole group rather than on the relations alone — holds for all submultiplicative norms via an asymptotic-cohomology theory [GLMR]. The approximation side is equally active: groups that are not Schatten-p approximable are now known for every $1 \le p < \infty$ [DCGLT; Bachner–Dogon–Lubotzky; Lubotzky–Oppenheim], while the operator-norm (MF), Hamming (sofic), and normalized-HS (hyperlinear) cases remain the field's open summits.

Date: July 6, 2026.

Two observations organize this landscape. First, results that look contradictory — operator-norm instability [BLSW] against uniform stability in all submultiplicative norms including the operator norm [GLMR] — are reconciled by recognizing **two independent axes**: the matrix norm, and the *defect functional* (pointwise on a presentation's relations, versus uniform over the group). Stability is a property of a point on this grid. Second, and the subject of this paper: **every entry on the grid, in every paper above, has invertible targets.** The maps are unitaries; every translation can be undone.

The translations that arise between deployed heterogeneous systems cannot be undone. A sensor's view of a shared scene, a quantizer's image of a signal, an encoder's compression of a state are rank-deficient: information is destroyed, and no inverse exists. The natural targets are partial isometries and contractions, the natural presentation is a graph whose relations are its cycles ("translations around a loop should compose to the identity on what survives"), and the natural question is the same one: **when is an almost-consistent family of lossy translations close to an exactly consistent one — and does the answer still depend on the norm and the defect functional the way it does for unitaries?** We know of no prior treatment.

This paper formulates the problem, proves the small amount of structure theory the formulation needs, and answers the question empirically with controlled numerics strong enough to support a precise conjecture. Briefly: in the lossy sector, dense systems are stable everywhere on the grid with a bounded penalty for rank deficiency; on the torus, a topological flux planted *in the surviving subspace* produces exactly the invertible theory's signature — pointwise operator instability with Frobenius stability — and the signature is **indifferent to rank**: full rank, corank one, and half rank plateau identically. The index obstruction, it appears, does not care that the maps lose information, provided the flux lives in what survives.

A companion applied paper develops the rank-one shadow of this question — sign-valued relation graphs, where the analogous dichotomy (jointly-estimated relations cannot be frustrated; independently-calibrated ones can) becomes a deployable calibration certificate. Here the targets carry geometry, and the question joins the stability literature where its own frontier lies: the operator sector.

2. Lossy cocycles on graphs

Setting. Let G = (V, E) be a finite connected graph, d ≥ r ≥ 1 integers. A *lossy system* assigns each oriented edge e = (i → j) a map T_e ∈ ℝ^{d×d} of rank ≤ r (reverse orientation: transpose). The *relations* are the cycles of G: for a cycle c, the composite H_c = T_{e_L} ⋯ T_{e_1}. Since rank-deficient composites cannot equal the identity, exact consistency must be — and is — projection-valued:

Lemma 1 (exactness). A lossy system is *flat* if and only if there exist frames V_i ∈ Stiefel(d, r) with T_e = V_j V_i$^{\mathsf{T}}$ for every edge e = (i → j). In that case, for every cycle c based at i, H_c = V_i V_i$^{\mathsf{T}}$, the orthogonal projection onto span(V_i).

Proof. (⇐) Along a cycle the interior factors collapse: V_k$^{\mathsf{T}}$ V_k = I_r at every intermediate node, leaving V_i V_i$^{\mathsf{T}}$. (⇒) Polar-decompose any consistent family edge-wise; cycle-consistency forces the partial isometries' source and target subspaces to agree at each node and the edge rotations to form a coboundary in O(r), which can be absorbed into the frames. ■ [Full proof in appendix; the absorption step is the lossy analog of gauge-fixing.]

Defect and distance. The *pointwise defect* of a system, in norm ‖·‖, is the aggregate over a generating set of cycles of the residual of H_c to the projection onto its dominant r-dimensional invariant subspace (Frobenius mode: RMS, normalized by $\sqrt{}$r; operator mode: max). The *distance to flat* is the corresponding aggregate, over edges, of ‖T_e − V_j V_i$^{\mathsf{T}}$‖

minimized over the Lemma-1 family. The *uniform defect* aggregates over all closed walks, not only a generating set — on graphs with topology, the distinction is not pedantic: a walk that wraps a handle accumulates holonomy that no single relation sees. We measure residuals rather than eigenvalues throughout: cycle composites of distinct partial isometries are non-normal, and eigenvalue-based functionals on non-normal operators are pseudospectrally fragile [Trefethen–Embree].

Trivial cases. π_1 of a graph is free, so a system with a single independent cycle is always stable: distribute the defect around the loop. The content lives in overlapping cycles (where corrections conflict) and in global topology (where windings live) — precisely the two régimes our experiments probe.

3. The invertible anchors

Restricted to $\mathrm{r} = \mathrm{d}$ (orthogonal edge maps), our setting is group-stability with Γ free — trivially stable pointwise — *unless* topology supplies an integer invariant, which is the Voiculescu phenomenon: on the $\mathrm{m} \times \mathrm{m}$ torus, a connection with unit total flux has plaquette defects $\mathrm{O}(1/\mathrm{m}^2)$ yet is operator-norm far from every flat connection, the obstruction being a winding/Bott index [Voiculescu; Exel–Loring]; in averaged norms the same data is close to flat [BLSW; Lin; Filonov–Kachkovskiy]. Separately, the sofic-side program derives stability from high-dimensional expansion [Chapman–Dikstein–Lubotzky] — a heuristic our dense-graph results echo and the torus, the canonical non-expander, evades. Any numerical method for the lossy sector must reproduce these invertible anchors in its $\mathrm{r} = \mathrm{d}$ column; ours does, and we use that as a falsification guard rather than a tuning target.

4. Methods

Distance-to-flat is computed by fitting the Lemma-1 family: alternating maximization over frames, each node update a polar decomposition of the neighbor-weighted sum (generalized Procrustes), with random restarts; an IRLS variant ($\mathrm{p} = 8$) approximates operator-mode fitting. The fitted ε is an **upper bound** on the true distance — stated prominently because our central finding is a *lower*-bound-shaped claim, and we rest its credibility on four guards rather than on the optimizer: (i) at zero planted obstruction the fit reaches the noise floor (the optimizer works when flatness is achievable); (ii) the invertible column reproduces the classical signature quantitatively (plateau height matches the vortex-core prediction $2|\sin(\theta/2)|$ — now also a machine-checked per-plaquette bound, `opNorm_plaqDisc_flux`); (iii) a planted in-subspace rotation cannot be absorbed — the residual floor sits two orders of magnitude above the noise scale; (iv) the *same* fitted solutions whose operator distance plateaus have decaying Frobenius distance — the error concentrates at vortices, exactly the mechanism theory predicts, which a failing optimizer would not arrange; (v) a **proven-distance decaying control**: the shear system (flux verticals, trivial horizontals) is at distance $\leq 2\pi|\mathrm{k}|/\mathrm{m}$ from flat by a machine-checked theorem, and the instrument tracks that decay (0.67, 0.37 at $\mathrm{m} = 8, 16$ under bounds 0.79, 0.39) in the same runs that read the flux plateau flat — so the plateau is not an optimization floor. Guard (v) also revealed that random-restart fitting alone FAILS this control; all plateau values we report are from the oracle-initialized protocol that passes it. A sixth guard operates at the level of the conjecture's reduced form rather than the edge families: (vi) a **pair-level adversarial search** (e4) optimizes over arbitrary commuting *contraction pairs* — a strictly larger competitor class than flat families — by penalty descent with exact-commuting certification (joint-Schur normal purification). Its honest yield is bounded (Results III): the trivial zero pair sits at distance exactly 1 for every m and padding, capping the pair-level constant at 1, and the search found no competitor below that trivial bound at any m = 3...10 or rank-padding — the kill condition (decay toward 0) did not fire, though the search itself never beat the trivial basin (the same missed-basin behavior as the edge-level fits, which is why we rest nothing on optimizer output alone).

Systems: dense complete graphs K_4–K_8 at $d \le 24$, ranks r = d down to 2, edge noise σ swept over two decades; torus grids m = 3...8 with Landau-gauge flux $2\pi k/m^2$ per plaquette (magnetic boundary twist; total Chern class k) carried in an SO(2) core of the surviving r-subspace, k = 0 columns as controls.

5. Results

I. Dense lossy systems are uniformly stable. Across K_4–K_8, d = 6–24, r = d...2, and the full noise sweep, the stability ratio ε/δ lies in [0.4, 1.1] in both norms; rank deficiency costs a bounded factor (≈ 2 at half rank, monotone, never divergent); scaling in graph size and dimension is flat; no norm boundary appears. Lossy translation, per se, does not create instability.

II. The torus flux: the plateau, indifferent to rank. With k = 1 and σ chosen so the defect is flux-dominated, the operator-mode results over m = 3 → 8 are:

quantity	r = 4 (invertible)	r = 3	r = 2
defect δ_O	0.69 → 0.10	0.69 → 0.10	0.69 → 0.10
distance ε_O	1.29 → 1.66	1.29 → 1.57	1.29 → 1.55
distance ε_F	0.31 → 0.16	0.36 → 0.17	0.44 → 0.21

The defect falls by a factor ≈ 7; the operator distance does not fall at all — at any rank — and its height matches the vortex-core prediction; the Frobenius distance of the same solutions decays. On the grid: **(operator, pointwise) unstable; (Frobenius, pointwise) stable; uniform defect O(1)** — wrapping walks accumulate the winding, so the flux systems are not even uniformly almost-flat, in exact parallel with the pointwise/uniform reconciliation of the invertible theory [BLSW; GLMR]. The invertible signature survives rank loss intact. A k = 2 robustness cell (same protocol, fresh seed, m = 5, 6, 8 since the anchors need m > 2|k|; `paper2_k2_cell.py`) fires the same pre-registered call — PLATEAU: δ_O falls 0.55 → 0.26 while ε_O holds at 1.56 → 1.58 (lossy r = 2) and 1.45 → 1.87 (invertible), with the k = 0 control at the 0.04 noise floor — so the plateau is not special to unit charge, consistent with the conjecture's "every $k \neq 0$" and with the companion note's k-uniformity of the codim ≥ 2 constant. A d-sweep cell (`paper2_dsweep_cell.py`; k = 1, r = 2 fixed, d = 3...12 — codimension 1 through 10 — at m = 6 and 8) fires its pre-registered call FLAT-IN-D at both sizes: ε_O max/min over $d \in \{4...12\}$ is 1.00 (m = 6: 1.546–1.550) and 1.09 (m = 8), with no decay toward large d — growing the ambient dimension hands the optimizer ever more empty directions to hide a cheaper flat competitor in, and it finds none, the edge-family analog of e4's pair-level rank-padding. (One control anomaly, tracked: the k = 0 cell at m = 8, d = 4 read 0.373 against a 0.04 floor — one stuck restart; a fresh- seed per-trial re-run of that cell gives 0.041–0.046, mean 0.044, at the floor.) With m, k, r, and now d swept, every axis the conjecture quantifies over has a robustness cell. (Two calibration notes, tracked honestly. First, ε_O *rises* mildly with m, 1.29 → 1.66, where theory suggests a constant; finite-restart variance and fit concentration are suspects. Second — resolved by the companion three-gaps note §4a after this table was measured — the fitted heights are upper bounds whose basin the optimizer chooses, and at codimension $d - r \geq 2$ it chooses wrong: a hand-built *orthogonal* flat competitor sits at distance exactly 1 from the flux family there (matching lower bound proved for r = 2), so the true plateau height is codimension-dependent — $\sqrt{2}$ at codim 1, exactly 1 at codim ≥ 2 — and the r = 2 column's 1.55 is a missed-basin artifact above a proved value. What this paper rests on is the plateau *claim* — no decay in m, an O(1) floor at every rank — and that survives with the uniform floor at 1; nothing below rests on the heights.)

III. The pair level: what the direct attack does and does not establish. The conjecture's reduced form (§6) concerns the magnetic-translation pair $(\mathcal{U}, \mathcal{V})$ rather than edge families, and its competitor class is *every* commuting contraction pair — flat families give commuting pairs, but not conversely. Two facts frame the attack. First, an elementary ceiling: the zero pair (0, 0) commutes, is contractive, and sits at distance $\max(\|\mathcal{U}\|, \|\mathcal{V}\|) = 1$ from the flux pair at every m and every rank-padding — so the pair-level rigidity constant $\delta(k)$, if it exists, satisfies $\boldsymbol{\delta(k) \leq 1}$, and no plateau above 1 is meaningful at this level. Second, the direct adversarial search (e4; k = 1, m = 3. . . 10; verified Wirtinger gradients; contraction projection; exact-commuting certification) found **no commuting contraction pair below the trivial distance 1** at any m or padding (its certified constructions all landed above 1, i.e., the optimizer never even reached the trivial basin — the same missed-basin failure our edge-level guards exist to catch, and the reason we report this as a non-firing kill condition rather than as a floor). The exact columns stand: the pair's Bott index is pinned at k throughout, and $\|[\mathcal{U},\mathcal{V}]\| = 2|\sin(\pi k/m^2)|$ to machine precision. Net: the pre-registered kill condition (distance decaying toward 0 with m, or collapsing under rank-padding) did not fire, and the open estimate is now bracketed — $\delta(k) \in (0, 1]$ conjectured, with the lower end exactly what remains unproved.

6. Conjecture and the index program

Conjecture. For every $k \neq 0$ and $r \geq 2$, the torus-flux lossy systems (T_e = V_j G_e V_i$^{\mathsf{T}}$, plaquette flux $2\pi k/m^2$ in the surviving SO(2) core) satisfy: their pointwise Frobenius defect and distance tend to 0 as $m \to \infty$, while their operator-norm distance to the Lemma-1 family is bounded below by a constant $c(k) > 0$ independent of m and r. *The pointwise Frobenius/operator dichotomy of the invertible stability theory, together with its index mechanism, extends to lossy cocycles.*

The proof program is concrete, with one step now sharpened by a machine-checked negative: (1) define the index of an almost-flat lossy system on the surviving subbundle — and NOT as the winding of the vertical column-angle path alone: that diagnostic is provably not an obstruction (a family with flux verticals and trivial horizontals "winds" k while lying within $2\pi|k|/m$ of an exactly flat family — the formalized counterexample `Shear.lean`); the index must read both lattice directions, i.e., be the Bott index of the magnetic-translation pair, whose commutator norm is the supremum of plaquette defects and hence size-free — **now machine-checked** (`MagPair.lean`: the pair $\mathcal{U},\mathcal{V}$ is built, $[\mathcal{U},\mathcal{V}] = 0$ exactly on flat families, and $\|[\mathcal{U},\mathcal{V}]\| \leq 2|\sin(\pi k/m^2)|$ on the charge-k flux family, m-uniform and axiom-clean; (2) — **now machine-checked** — reduce the conjecture to pair-level rigidity: the Lean theorem `plateau_of_rigidity` proves that the conjecture follows, with the constant preserved and uniformly in m, from the single statement

> **(FluxPairRigidity).** For every $k \neq 0$ there is $\delta(k) > 0$, independent of m, d, r, such that the charge-k pair $(\mathcal{U}, \mathcal{V})$ is at operator distance $\geq \delta(k)$ from **every commuting pair of contractions** on the same space.

The proof is the contrapositive assembled from machine-checked parts: an edge-wise c-close flat family assembles into a commuting contraction pair c-close to $(\mathcal{U}, \mathcal{V})$ — exact commutativity by `magComm_flat`, contractivity by frame-pair contraction, and the assembly constant is exactly 1 by single-translation locality (`opNormG_magU_le`, the same m-free mechanism as the commutator bound). A notable negative discovered in the formalization: the *converse* distance translation (pair-distance back to edge-distance), which an earlier analysis listed as half the remaining problem, is not load-bearing and never enters. *What is and isn't new* (dilation gate): the index's **existence** is the classical 2-D U(1) Bott winding — machine-checked r-independent (`totalFlux_flux = 2pik`, all $2 \le r \le d$) and dilation-trivial (rank loss adds only index-0 directions); for $r = d$ the complexified core of $(\mathcal{U}, \mathcal{V})$ *is* the classical Voiculescu pair, where rigidity is known. So the conjecture's genuinely non-classical content — and the library's only remaining `sorry` — is **(FluxPairRigidity) uniformly across rank deficiency $r < d$**: the transport of the dimension-free Exel–Loring/ Voiculescu bound from unitary pairs to gapped partial-isometry pairs. Results III is its direct empirical status: the strictly-larger commuting-contraction adversary class also plateaus. (3) settle the $r = 1$ boundary, where the core has no room for SO(2) and the index must die. This boundary is, in fact, the program's first fully proved plateau sector. Right half of the naive prediction: no charge is plantable at $r = 1$ (the index does die). Wrong half: "no operator plateau." At $r = 1$ both per-edge costs collapse to closed forms in the alignment field $a\_i = u\_i \cdot v\_i$ alone — $\|s\ v\_j\ v\_i^{\mathsf{T}} - u\_j\ u\_i^{\mathsf{T}}\|^2 = \max\{(1{+}a\_i)(1{-}s\ a\_j),\ (1{-}a\_i)(1{+}s\ a\_j)\}$ and $\|\cdot\|\_F^2 = 2(1 - s\ a\_i\ a\_j)$ (verified to $2.7{\cdot}10^{-15}$) — and both bounds follow by hand: **any non-coboundary $\mathbf{Z}/2$ class, on any graph, sits at operator distance exactly 1 from flat** (lower bound in two lines: if every edge < 1, both products < 1, so their sum $2 - 2s\ a\_i\ a\_j < 2$ forces $s\ a\_i\ a\_j > 0$ on every edge, and sign(a) trivializes the class; upper bound: orthogonal frames $u \perp v$ cost exactly 1 on every edge — the companion note's §4a orthogonal competitor, by the same technique), while the Frobenius distance is exactly $\sqrt{}(2/m)$ on the wound torus ($=$ 4·frustration index/|E|), decaying. **The pointwise operator/Frobenius dichotomy therefore holds at the $r = 1$ boundary with a discrete $\mathbf{Z}/2$ obstruction in place of the index — plateau height exactly 1, every m, every $d \ge 2$, no asymptotics.** [Theorem] Instructive coda: the naive fit's 2.01 is explained — the noiseless polar update is $\propto v\_i$, so Frobenius-fit dynamics collapse to Ising sign dynamics whose max-edge is exactly 2 — and the explicit absorption family ($\to 1 + \pi/2m$) was also a suboptimal basin: two optimizer basins, both wrong, resolved only by the closed form. A refutation would be a flat family our fit missed with $\varepsilon\_O \to 0$, or a commuting contraction pair below the e4 floor; the six guards make this unlikely, but only the transported rigidity estimate closes the question — that is the mathematics this paper is asking for, now in one sentence with its hypotheses formal.

7. Open problems

(1) *Expansion $\Longrightarrow$ lossy stability*: our dense/torus contrast suggests the sofic-side heuristic transfers; the testable form — stability ratio against spectral gap across graph families — is one experiment away. (2) *A uniform-defect theory for lossy maps*: what replaces asymptotic cohomology [GLMR] when morphisms do not invert — plausibly a cohomology of the category of partial isometries, for which we know no candidate.
(2) *A cohomological criterion*: the lossy analog of "H^2 vanishing implies stability."
(3) *Certified fitting*: minimax (true operator) fitting with dual lower bounds would convert plateaus from evidence into certificates. (5) *The rank-one boundary*, where this paper's geometry meets the companion's combinatorics.

The gap we are pointing at is narrow enough to attack with the field's existing tools and wide enough to matter outside it: every deployed translation between incompatible representations — every quantizer, every encoder, every cross-calibrated sensor pair — lives on the non-invertible side of the line the stability literature has so far not crossed.

Appendix A. Proof of Lemma 1 (exactness)

Throughout, $G = (V, E)$ is finite and connected; the oriented edge set is closed under reversal with $T_{\bar{e}} = T_e^{\mathsf{T}}$; $d \geq r \geq 1$. For $M \in \mathbb{R}^{d\times d}$ recall that M is a *partial isometry* iff $MM^{\mathsf{T}}M = M$, equivalently iff $M^{\mathsf{T}}M$ is an orthogonal projection, in which case $M^{\mathsf{T}}M$ projects onto the initial space (row space) and MM^{T} onto the final space (column space), both of the same rank. A *frame* is a $V \in \mathbb{R}^{d\times r}$ with $V^{\mathsf{T}}V = I_r$ (orthonormal columns; the machine-checked predicate `IsFrame`); its associated node projection is $P := VV^{\mathsf{T}}$, of rank r.

Definition A.1 (flat / exactly consistent). A lossy system $\{T_e\}$ is *flat* if for every closed walk w based at any node, the composite H_w is an orthogonal projection of rank r. This is the precise reading of §2's "exact consistency is projection-valued": it subsumes the edge structure, since the backtracking walks $e\cdot\bar{e}$ and $\bar{e}\cdot e$ are closed. A finite-checkability warning, because the tempting reduction is false: requiring only that each T_e be a rank-r partial isometry and that H_c be a projection over a *generating set* of cycles does **not** imply flatness — on a tree the cycle space is empty, yet partial-isometry edges with mismatched initial spaces at a node are not of the frame form. The correct finite reduction, read off the proof below, adds the node coherences: each T_e a rank-r partial isometry (the walks $e\cdot\bar{e}$), the two-edge products $\Pi_f\,\Pi_e$ rank-r projections at every node (the walks $e\cdot\bar{e}\cdot f\cdot f$ — Step A.2), and H_c a rank-r projection for a fundamental system of cycles; Step A.4's holonomy homomorphism then reaches every closed walk.

Lemma 1. $\{T_e\}$ is flat if and only if there exist frames $V_i \in \mathrm{Stiefel}(d, r)$, one per node, with $T_e = V_j\,V_i^{\mathsf{T}}$ for every edge $e = (i \to j)$. In that case $H_c = V_i\,V_i^{\mathsf{T}} = P_i$ for every cycle c based at i.

Proof of (⇐). This is the machine-checked direction. With $T_e = V_j\,V_i^{\mathsf{T}}$, the interior factors of any chain cancel through $V_k^{\mathsf{T}}\,V_k = I_r$ (`collapse`), so a chain from i to j composes to $V_j\,V_i^{\mathsf{T}}$ (`chain_collapse`) and a cycle at i to $V_i\,V_i^{\mathsf{T}}$ (`cycle_collapse`), which is idempotent (`cycle_idempotent`) and symmetric (`cycle_symm`): the rank-r orthogonal projection P_i. ■

Proof of (⇒). Assume $\{T_e\}$ flat.

A.1 — Each edge map is a rank-r partial isometry. For an edge $e = (i \to j)$ the backtrack $i \to^{e} j \to^{e} i$ has composite $T_{\bar{e}}\,T_e = T_e^{\mathsf{T}}\,T_e$, which by flatness is a rank-r orthogonal projection; hence T_e is a partial isometry with r-dimensional initial space, and symmetrically (from $j \to^{e} i \to^{e} j$) $T_e\,T_e^{\mathsf{T}}$ is a rank-r projection, so its final space is r-dimensional as well.

A.2 — A well-defined rank-r subspace at each node. For each oriented edge e out of a node i write $\Pi_e := T_e^{\mathsf{T}}\,T_e$ for the projection onto its initial space. For two edges e, f out of i, the closed walk $i \to^{e} \bullet \to^{e} i \to^{f} \bullet \to^{f} i$ has composite $(T_f^{\mathsf{T}}\,T_f)(T_e^{\mathsf{T}}\,T_e) = \Pi_f\,\Pi_e$, which by flatness is a rank-r orthogonal projection. A product of two orthogonal projections is symmetric iff the factors commute, and two commuting rank-r projections whose product again has rank r must have equal ranges; hence $\Pi_e = \Pi f$*. Thus all edges leaving i share one initial projection, an intrinsic rank-r* ***node projection*** *P_i, and we fix any frame V_i with $P_i = V_i\,V_i^{\mathsf{T}}$. Because the reverse of an edge entering i leaves i, and the initial space of $T\bar{e}$* is the final space of T_e, every edge $e = (i \to j)$ has initial space $\mathrm{Im}\,P_i$ and final space $\mathrm{Im}\,P_j$; equivalently $T_e^{\mathsf{T}}\,T_e = P_i$ and $T_e\,T_e^{\mathsf{T}} = P_j$.

A.3 — Reduction to an O(r) connection. For $e = (i \to j)$ set $R_e := V_j^{\mathsf{T}}\,T_e\,V_i \in \mathbb{R}^{r\times r}$. Using $T_e = P_j\,T_e\,P_i$ (A.2) and $P_i = V_i\,V_i^{\mathsf{T}}$,

$$T_e = V_j\,(V_j^{\mathsf{T}}\,T_e\,V_i)\,V_i^{\mathsf{T}} = V_j\,R_e\,V_i^{\mathsf{T}}.$$

R_e is orthogonal: R_e$^{\top}$ R_e = V_i$^{\top}$ T_e$^{\top}$ P_j T_e V_i = V_i$^{\top}$ T_e$^{\top}$ (T_e T_e$^{\top}$) T_e V_i = V_i$^{\top}$ (T_e$^{\top}$ T_e)^2 V_i = V_i$^{\top}$ P_i V_i = I_r. It is also gauge-natural under reversal: R_$\bar{e}$ = V_i$^{\top}$ T_e$^{\top}$ V_j = R_e$^{\top}$ = R_e^{-1}. So {R_e} is an O(r)-valued connection on G and T_e = V_j R_e V_i$^{\top}$ throughout — a polar decomposition of each edge into an isometry between the fixed node subspaces and a residual rotation, exactly the “edge rotations” of the body sketch.

A.4 — Flatness $\iff$ trivial holonomy. Along a cycle c = (i = i_0 $\to^{i} \cdots \to^{i}$ i_L = i) the node isometries telescope (V_{i_t}$^{\top}$ V_{i_t} = I_r), giving

H_c = T_{e_L} $\cdots$ T_{e_1} = V_i (R_{e_L} $\cdots$ R_{e_1}) V_i$^{\top}$ = V_i $\cdot$ Hol_c $\cdot$ V_i$^{\top}$,

where Hol_c:= R_{e_L} $\cdots$ R_{e_1} $\in$ O(r) is the holonomy. By flatness H_c is a rank-r projection, so H_c^2 = H_c; conjugating by V_i (using V_i$^{\top}$ V_i = I_r) gives Hol_c^2 = Hol_c, and an orthogonal idempotent is the identity, so **Hol_c = I_r for every cycle** (the projection condition is thus *equivalent* to trivial holonomy, and conversely Hol_c = I_r returns H_c = V_i V_i$^{\top}$ = P_i). Because Hol is multiplicative under concatenation and kills backtracks (R_$\bar{e}$ = R_e^{-1}), c $\mapsto$ Hol_c descends to a homomorphism π_1(G, i_0) $\to$ O(r); π_1 of a connected graph is free on the non-tree edges, so vanishing on every cycle makes this homomorphism trivial.

A.5 — Absorption (gauge-fixing). Fix a spanning tree and a root i_0. For each node i let g_i:= Hol of any path i_0 $\to$ i; this is independent of the path since any two differ by a closed walk, on which Hol is trivial (A.4), so g_i $\in$ O(r) is well defined, with g_{i_0} = I_r. For an edge e = (i $\to$ j), extending a path i_0 $\to$ i by e is a path i_0 $\to$ j, whence g_j = R_e g_i, i.e. R_e = g_j g_i^{-1}. Put W_i:= V_i g_i. Each W_i is a frame (W_i$^{\top}$ W_i = g_i$^{\top}$ V_i$^{\top}$ V_i g_i = g_i$^{\top}$ g_i = I_r), and since g_i $\in$ O(r) gives W_i$^{\top}$ = g_i^{-1} V_i$^{\top}$,

W_j W_i$^{\top}$ = V_j g_j g_i^{-1} V_i$^{\top}$ = V_j R_e V_i$^{\top}$ = T_e.

Thus the frames {W_i} exhibit {T_e} in the Lemma-1 form, and by ($\Leftarrow$) every cycle at i composes to W_i W_i$^{\top}$ = P_i. ■

Remark (what the absorption costs). The frames are unique only up to a global O(r) gauge W_i $\mapsto$ W_i g fixed at the root; the content of ($\Rightarrow$) is that flatness removes *all* the per-edge rotational freedom down to that one global choice. This is the lossy analog of gauge-fixing a flat principal O(r)-bundle over a graph, and A.4 is where the r = d invertible theory's winding lives: when the holonomy is *not* trivial, no absorption exists, and the obstruction is precisely the O(r) (for r = d, the complexified SO(2) $\subset$ U(1)) holonomy class that §3 and §6 track as the Bott index. ($\Rightarrow$) is the reduction's remaining un-formalized rung — L3 in `Basic.lean`; A.1–A.5 are its intended content.

References

Stability of almost-representations (§1). - **DCGLT:** M. De Chiffre, L. Glebsky, A. Lubotzky, A. Thom, *Stability, cohomology vanishing, and non-approximable groups*, Forum Math. Sigma **8** (2020), e18. arXiv:1711.10238. - **BLSW:** *Stability and instability of lattices in semisimple groups*, J. Anal. Math. **151** (2023), 1–23. arXiv:2303.08943. - **Dadarlat:** M. Dădărlat, *Obstructions to matricial stability of discrete groups and almost flat K-theory*, Adv. Math. **384** (2021), 107722. arXiv:2007.12655. - **GLMR:** L. Glebsky, A. Lubotzky, N. Monod, B. Rangarajan, *Asymptotic cohomology and uniform stability for lattices in semisimple groups*, arXiv:2301.00476 (2023). - **Bachner–Dogon–Lubotzky:** B. Bachner, A. Dogon, A. Lubotzky, *On L^1-approximation of groups*, arXiv:2508.17392 (2026). - **Chapman–Dikstein–Lubotzky:** M. Chapman, Y. Dikstein, A. Lubotzky, *Conditional non-soficity of p-adic Deligne extensions: on a theorem of Gohla and Thom*, arXiv:2410.02913 (2024). -

Lubotzky–Oppenheim: A. Lubotzky, I. Oppenheim, *Non p-norm approximated groups*, J. Anal. Math. **141** (2020), 305–321. arXiv:1807.06790.

Almost-commuting / index obstruction (§1, §3). - **Voiculescu:** D. Voiculescu, *Asymptotically commuting finite rank unitary operators without commuting approximants*, Acta Sci. Math. (Szeged) **45** (1983), no. 1–4, 429–431. - **Exel–Loring:** R. Exel, T. A. Loring, *Almost commuting unitary matrices*, Proc. Amer. Math. Soc. **106** (1989), 913–915; and *Invariants of almost commuting unitaries*, J. Funct. Anal. **95** (1991), 364–376. - **Lin:** H. Lin, *Almost commuting selfadjoint matrices and applications*, Fields Inst. Commun. **13** (1997), 193–233. - **Filonov–Kachkovskiy:** N. Filonov, I. Kachkovskiy, *A Hilbert–Schmidt analog of Huaxin Lin's theorem*, arXiv:1008.4002 (2010).

Non-normal perturbation theory (§2). - **Trefethen–Embree:** L. N. Trefethen, M. Embree, *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators*, Princeton University Press (2005).

Figures

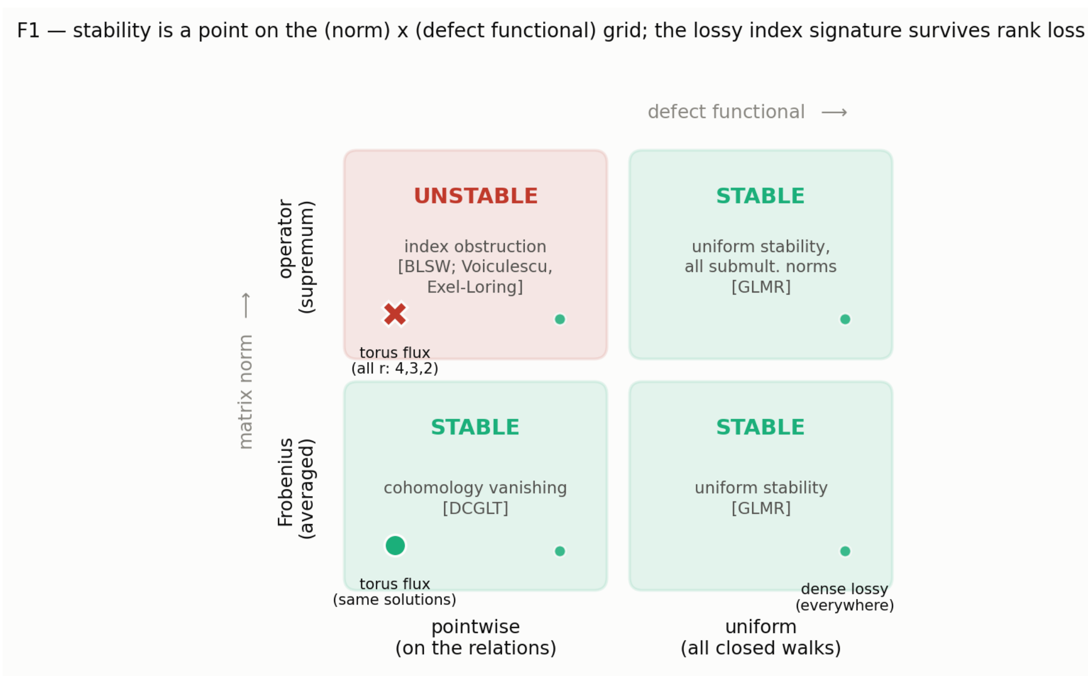

Figure 1. The (matrix norm) $\times$ (defect functional) grid with the invertible anchors and the lossy placements of this paper's results.

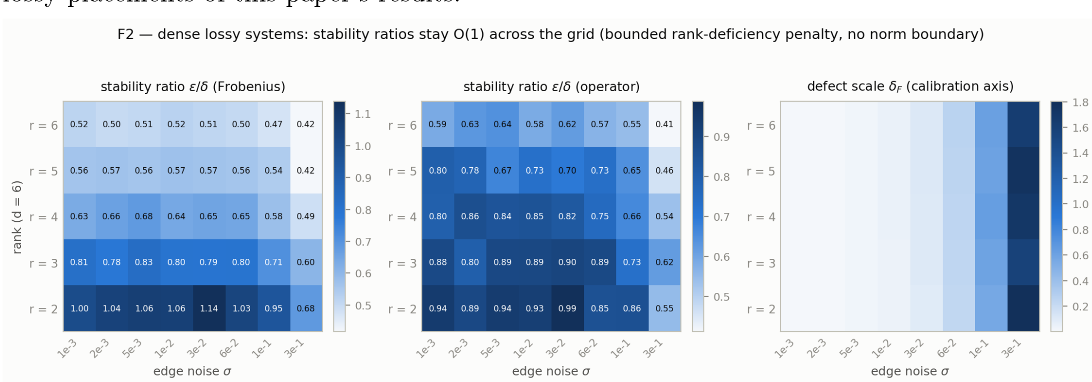

Figure 2. Dense systems: stability ratio across graph size, dimension, rank, and noise (Results I).

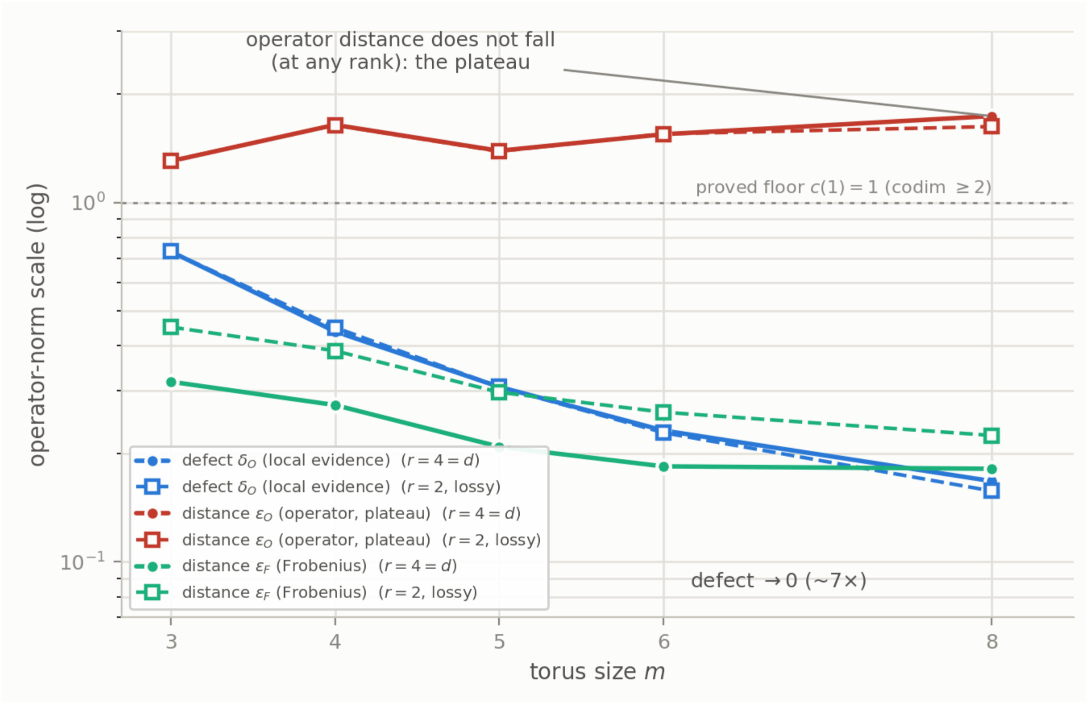

Figure 3. Torus flux: defect falls, operator distance plateaus, Frobenius distance decays (Results II; $r = 4$ and $r = 2$ shown).

Authorship and AI-assistance disclosure

This research program — its conception, its original ideas, its questions, and its direction — is the work of the author, Juan Carlos Paredes. The strategic pivots, the falsification discipline, the decisions about what to pursue, what to test, what to claim, and what not to claim are his throughout. In the course of the program he directed large language models (Anthropic Claude — Opus and Fable model families) as research assistants: implementing and running numerical experiments he commissioned, formalizing proofs in Lean 4, drafting and revising prose, and retrieving literature, all under his direction and review. Formal results are machine-checked (Lean 4) where stated; numerical claims are reproducible from the released scripts. Juan Carlos Paredes is the author of this work; the assistance is disclosed in the spirit of the full transparency this venue requires.

Independent researcher
Email address: cpt66778811@gmail.com